

Ontology-Based Access to Probabilistic Data

Jean Christoph Jung and Carsten Lutz

Universität Bremen, Germany
{jeanjung,clu}@informatik.uni-bremen.de

Abstract. We propose a framework for querying probabilistic data in the presence of an ontology, arguing that the interplay of probabilities and ontologies is fruitful in applications such as managing data that was extracted from the web. The prime inference problem is computing answer probabilities, and we show that it can be implemented using standard probabilistic database systems, similar to traditional ontology-based data access. We demonstrate that query rewriting into first-order logic is an important tool for our framework. First, it is used to establish a PTIME vs. #P dichotomy for the data complexity of this problem by lifting a corresponding result from probabilistic databases. Then, we use it to characterize which pairs of query and TBox are in PTIME. Finally, it is shown that non-existence of such a rewriting implies #P-hardness.

1 Introduction

In recent years, *ontology based data access* (OBDA) has become an active area of description logic research. In OBDA, an ontology provides a semantics for *incomplete* data with the aim of facilitating the computation of more complete answers to queries. There are applications, though, in which it is necessary to query data that is not only incomplete, but also *uncertain*. For instance, data extracted from web sources [24] such as an estate agents' web page is typically incomplete. It is also uncertain because web sources tend to be unreliable and extraction tools are based on heuristic decisions and thus significantly error prone. In this paper, we propose an extension of OBDA that captures uncertain data through a probabilistic data model and replaces the computation of certain answers with computing the probabilities of certain answers. In brief, our approach relates to probabilistic database systems (PDBMSs) in the same way that traditional OBDA relates to RDBMSs.

Framework. We consider probabilistic ABox formalisms that are inspired by data models from the currently very active area of probabilistic databases [7,32]. Specifically, *pABoxes* enrich classical ABoxes with probabilities that are attached to *probabilistic events*, and with *event expressions* that are attached to ABox assertions. For example, a pABox assertion `SoccerPlayer(messi)` can be associated with an event expression $e_1 \vee e_2$, where e_1 and e_2 represent events such as 'web extraction tool x correctly analyzed webpage y stating that Messi is a soccer player'. Event e_1 can then be associated with probability .7, and e_2 with .9. Events are assumed to be probabilistically independent, which results

in a straightforward semantics that is similar to well-known probabilistic versions of datalog [28,12]. Ontologies are represented by description logic TBoxes. In this setting, which we call *ontology based access to probabilistic data (pOBDA)*, we are interested in *computing the probabilities of certain answers to conjunctive queries (CQs)*. Note that uncertainty occurs only in the data, but neither in the ontology nor in the query. We believe that pOBDA is of general interest and potentially useful for a wide range of applications including the management of data extracted from the web, machine translation, and dealing with data that arises from sensor networks. All these applications can potentially benefit from a fruitful interplay between ontologies and probabilities; in particular, the ontology can help to reduce the uncertainty of the data.

Contributions. The main aim of this paper is to study the *data complexity* of pOBDA. More precisely, we pursue a *non-uniform* approach as recently initiated in [27], which aims at fully classifying the data complexity of every pair (q, \mathcal{T}) that consists of a CQ q and a TBox formulated in a fixed ‘master DL’. As a central tool of our analysis, we use *query rewriting* into first-order (FO) queries, which is an important technique for traditional OBDA [6]. We start with showing that FO-rewritings from traditional OBDA are useful also in the context of pOBDA: for any pABox \mathcal{A} , the probability that a tuple \mathbf{a} is a certain answer to q over a pABox \mathcal{A} relative to \mathcal{T} is identical to the probability that \mathbf{a} is an answer to the FO-rewriting $q_{\mathcal{T}}$ of q and \mathcal{T} over \mathcal{A} viewed as a probabilistic database. This *lifting* of FO-rewritings to the probabilistic case immediately implies that one can implement pOBDA based on existing PDBMSs such as MayBMS, Trio, and MystiQ [1,35,5].

Lifting also allows us to carry over the dichotomy between PTIME and #P-hardness for computing the probabilities of answers to unions of conjunctive queries (UCQs) over probabilistic databases recently obtained by Dalvi, Suciu, and Schnaitter [8] to our pOBDA framework, provided that we restrict ourselves to TBoxes formulated in (the core version of) DL-Lite and to ipABoxes, which are a special case of pABoxes in which all ABox assertions are probabilistically independent. Based on a careful syntactic analysis, we provide a concrete characterization of those CQs q and DL-Lite TBoxes \mathcal{T} for which computing answer probabilities is in PTIME. We then proceed to showing that query rewriting is a *complete* tool for proving PTIME data complexity in pOBDA, in the following sense: we replace DL-Lite with the strictly more expressive description logic \mathcal{ELI} where, in contrast to DL-Lite, rewritings into first-order queries do not exist for every CQ q and TBox \mathcal{T} ; we then prove that if some (q, \mathcal{T}) does *not* have a rewriting, then computing answer probabilities for q relative to \mathcal{T} is #P-hard. Thus, if it is possible at all to prove that some (q, \mathcal{T}) has PTIME data complexity, then this can always be done using query rewriting. Both in DL-Lite and \mathcal{ELI} , the class of queries and TBoxes with PTIME data complexity is relatively small. This negative result is relativized by the fact that answer probabilities can often be efficiently approximated. In particular, all pairs (q, \mathcal{T}) admit approximation in terms of a *fully polynomial randomized approximation scheme (FPRAS)* whenever q is FO-rewritable relative to \mathcal{T} . We provide a brief

discussion of such approximations and refer to the full version [19] of this paper for more information.

Related Work. The probabilistic ABox formalism studied in this paper is inspired by the probabilistic database models in [9], but can also be viewed as a variation of probabilistic versions of datalog and Prolog, see [28,12] and references therein. There have recently been other approaches to combining ontologies and uncertainty for data access [11,14], with a different semantics; the setup considered by Gottlob, Lukasiewicz, and Simari in [14] is close in spirit to the framework studied here, but also allows probabilities in the TBox and has a different, rather intricate semantics based on Markov logic. In fact, we deliberately avoid probabilities in the ontology because (i) this results in a simple and fundamental, yet useful formalism that still admits a very transparent semantics and (ii) it enables the use of standard PDBMSs for query answering in the presence of ontologies. Note that the analogous property of being implementable using state-of-the-art RDBMSs is a favorable feature of traditional OBDA [6]. There has also been a large number of proposals for enriching description logic TBoxes (instead of ABoxes) with probabilities, see [25,26] and the references therein. Our running application example is web data extraction, in the spirit of [16] to store extracted web data in a probabilistic database. Note that it has also been proposed to integrate both probabilities and ontologies directly into the data extraction tool [13]. We believe that both approaches can be useful and could even be orchestrated to play together.

This paper is a condensed version of [19]. For the missing proofs and some additional material we refer the reader to the long version of [19].

2 Preliminaries

We use standard notation for the syntax and semantics of description logics (DLs) and refer to [3] for full details. As usual, C, D denote (potentially) composite concepts, A, B concept names, r, s role names, R and S role names or their inverse, and a, b individual names. When $R = r^-$, then R^- denotes r . We consider the ontology languages $DL\text{-Lite}$ and \mathcal{ELI} . Regarding the former, we concentrate on the dialect $DL\text{-Lite}_{\text{core}}$, where $TBoxes$ are finite sets of *concept inclusions* (CIs) $B \sqsubseteq B'$ and $B \sqcap B' \sqsubseteq \perp$ with B and B' concepts of the form $\exists r, \exists r^-, \top$ or A . In \mathcal{ELI} , a $TBox$ is a finite set of CIs $C \sqsubseteq D$ where C and D are (potentially) compound concepts of the form $\top, A, C' \sqcap D', \exists r.C'$, and $\exists r^-.C'$.

As usual, an $ABox$ is a finite set of *concept assertions* $A(a)$ and *role assertions* $r(a, b)$. We use $\text{Ind}(\mathcal{A})$ to denote the set of individual names used in the $ABox$ \mathcal{A} and write $r^-(a, b) \in \mathcal{A}$ for $r(b, a) \in \mathcal{A}$. The semantics of DLs is based on interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$; we adopt the unique name assumption (UNA), i.e., we require $a^{\mathcal{I}} \neq b^{\mathcal{I}}$ for all individuals $a \neq b$.

Conjunctive queries (CQs) take the form $\exists \mathbf{y}.\varphi(\mathbf{x}, \mathbf{y})$, with φ a conjunction of atoms of the form $A(t)$ and $r(t, t')$ and where \mathbf{x}, \mathbf{y} denote (tuples of) variables and t, t' denote *terms*, i.e., a variable or an individual name. We call the variables

in \mathbf{x} the *answer variables* and those in \mathbf{y} the *quantified variables*. The set of all variables in a CQ q is denoted by $\text{var}(q)$ and the set of all terms in q by $\text{term}(q)$. A CQ q is *n-ary* if it has n answer variables and *Boolean* if it is 0-ary. Whenever convenient, we treat a CQ as a *set* of atoms and sometimes write $r^-(t, t') \in q$ instead of $r(t', t) \in q$.

Let \mathcal{I} be an interpretation and q a CQ with answer variables x_1, \dots, x_k . For individual names $\mathbf{a} = a_1 \cdots a_k$, an *\mathbf{a} -match* for q in \mathcal{I} is a mapping $\pi : \text{term}(q) \rightarrow \Delta^{\mathcal{I}}$ such that $\pi(x_i) = a_i$ for $1 \leq i \leq k$, $\pi(a) = a^{\mathcal{I}}$ for all individual names $a \in \text{term}(q)$, $\pi(t) \in A^{\mathcal{I}}$ for all $A(t) \in q$, and $(\pi(t_1), \pi(t_2)) \in r^{\mathcal{I}}$ for all $r(t_1, t_2) \in q$. We write $\mathcal{I} \models q[\mathbf{a}]$ if there is an \mathbf{a} -match of q in \mathcal{I} and let $\text{ans}(q, \mathcal{I})$ denote the set of all \mathbf{a} with $\mathcal{I} \models q[\mathbf{a}]$. For a TBox \mathcal{T} and an ABox \mathcal{A} , we write $\mathcal{T}, \mathcal{A} \models q[\mathbf{a}]$ if $\mathcal{I} \models q[\mathbf{a}]$ for all models \mathcal{I} of \mathcal{T} and \mathcal{A} . The set $\text{cert}_{\mathcal{T}}(q, \mathcal{A})$ of all *certain answers* consists of all tuples \mathbf{a} over $\text{Ind}(\mathcal{A})$ with $\mathcal{T}, \mathcal{A} \models q[\mathbf{a}]$.

3 Probabilistic OBDA

We introduce our framework for probabilistic OBDA, starting with the definition of a rather general, probabilistic version of ABoxes. Let \mathcal{E} be a countably infinite set of *atomic (probabilistic) events*. An *event expression* is built up from atomic events using the Boolean operators \neg, \wedge, \vee . We use $\text{expr}(\mathcal{E})$ to denote the set of all event expressions over \mathcal{E} . A *probability assignment for $E \subseteq \mathcal{E}$* is a map $E \rightarrow [0, 1]$.

Definition 1 (pABox). A probabilistic ABox (pABox) is of the form (\mathcal{A}, e, p) with \mathcal{A} an ABox, e a map $\mathcal{A} \rightarrow \text{expr}(\mathcal{E})$, and p a probability assignment for $E_{\mathcal{A}}$, the atomic events in \mathcal{A} .

Example 1. We consider as a running example an information extraction tool that is gathering data from the web, see [16] for a similar setup. Assume we are gathering data about soccer players and the clubs they play for in the current 2012 season, and we want to represent the result as a pABox.

(1) The tool processes a newspaper article stating that ‘Messi is the soul of the Argentinian national soccer team’. Because the exact meaning of this phrase is unclear (it could refer to a soccer player, a coach, a mascot), it generates the assertion $\text{Player}(\text{messi})$ associated with the atomic event expression e_1 with $p(e_1) = 0.7$. The event e_1 represents that the phrase was interpreted correctly.

(2) The tool finds the Wikipedia page on Lionel Messi, which states that he is a soccer player. Since Wikipedia is typically reliable and up to date, but not *always* correct, it updates the expression associated with $\text{Player}(\text{messi})$ to $e_1 \vee e_2$ and associates e_2 with $p(e_2) = 0.95$.

(3) The tool finds an HTML table on the homepage of FC Barcelona saying that the top scorers of the season are Messi, Villa, and Pedro. It is not stated whether the table refers to the 2011 or the 2012 season, and consequently we generate the ABox assertions $\text{playsfor}(x, \text{FCbarca})$ for $x \in \{\text{messi}, \text{villa}, \text{pedro}\}$ all associated with the same atomic event expression e_3 with $p(e_3) = 0.5$. Intuitively, the event e_3 is that the table refers to 2012.

- (4) Still processing the table, the tool applies the background knowledge that top scorers are typically strikers. It generates three assertions $\text{Striker}(x)$ with $x \in \{\text{messi, villa, pedro}\}$, associated with atomic events $e_4, e'_4,$ and e''_4 . It sets $p(e_4) = p(e'_4) = p(e''_4) = 0.8$. The probability is higher than in (3) since being a striker is a more stable property than playing for a certain club, thus this information does not depend so much on whether the table is from 2011 or 2012.
- (5) The tool processes the twitter message ‘Villa was the only one to score a goal in the match between Barca and Real’. It infers that Villa plays either for Barcelona or for Madrid, generating the assertions $\text{playsfor}(\text{villa, FCbarca})$ and $\text{playsfor}(\text{villa, realmadrid})$. The first assertion is associated with the event e_5 , the second one with $\neg e_5$. It sets $p(e_5) = 0.5$.

We now define the semantics of OBDA over pABoxes. Each $E \subseteq E_{\mathcal{A}}$ can be viewed as a truth assignment that makes all events in E true and all events in $E_{\mathcal{A}} \setminus E$ false.

Definition 2. Let (\mathcal{A}, e, p) be a pABox. For each $E \subseteq E_{\mathcal{A}}$, define a corresponding non-probabilistic ABox $\mathcal{A}_E := \{\alpha \in \mathcal{A} \mid E \models e(\alpha)\}$. The function p represents a probability distribution on $2^{E_{\mathcal{A}}}$, by setting for each $E \subseteq E_{\mathcal{A}}$:

$$p(E) = \prod_{e \in E} p(e) \cdot \prod_{e \in E_{\mathcal{A}} \setminus E} (1 - p(e)).$$

The probability of an answer $\mathbf{a} \in \text{Ind}(\mathcal{A})^n$ to an n -ary conjunctive query q over a pABox \mathcal{A} and TBox \mathcal{T} is

$$p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) = \sum_{E \subseteq E_{\mathcal{A}} : \mathbf{a} \in \text{cert}_{\mathcal{T}}(q, \mathcal{A}_E)} p(E).$$

For Boolean CQs q , we write $p(\mathcal{A}, \mathcal{T} \models q)$ instead of $p_{\mathcal{A}, \mathcal{T}}(\emptyset \in q)$, where (\emptyset) denotes the empty tuple.

Example 2. Consider again the web data extraction example discussed above. To illustrate how ontologies can help to reduce uncertainty, we use the DL-Lite TBox

$$\mathcal{T} = \left\{ \begin{array}{ll} \exists \text{playsfor} \sqsubseteq \text{Player} & \text{Player} \sqsubseteq \exists \text{playsfor} \\ \exists \text{playsfor}^- \sqsubseteq \text{SoccerClub} & \text{Striker} \sqsubseteq \text{Player} \end{array} \right\}$$

Consider the following subcases considered above.

- (1) + (3) The resulting pABox comprises the following assertions with associated event expressions:

$$\begin{array}{ll} \text{Player}(\text{messi}) \rightsquigarrow e_1 & \text{playsfor}(\text{messi, FCbarca}) \rightsquigarrow e_3 \\ \text{playsfor}(\text{villa, FCbarca}) \rightsquigarrow e_3 & \text{playsfor}(\text{pedro, FCbarca}) \rightsquigarrow e_3 \end{array}$$

with $p(e_1) = 0.7$ and $p(e_3) = 0.5$. Because of the statement $\exists \text{playsfor} \sqsubseteq \text{Player}$, using \mathcal{T} (instead of an empty TBox) increases the probability of messi to be an answer to the query $\text{Player}(x)$ from 0.7 to 0.85.

(5) The resulting pABox is

$$\text{playsfor}(\text{villa}, \text{FCbarca}) \rightsquigarrow e_5 \quad \text{playsfor}(\text{villa}, \text{realmadrid}) \rightsquigarrow \neg e_5$$

with $p(e_5) = 0.5$. Although $\text{Player}(\text{villa})$ does not occur in the data, the probability of villa to be an answer to the query $\text{Player}(x)$ is 1 (again by the TBox-statement $\exists \text{playsfor} \sqsubseteq \text{Player}$).

(3)+(4) This results in the pABox

$$\begin{aligned} \text{playsfor}(\text{messi}, \text{FCbarca}) &\rightsquigarrow e_3 & \text{Striker}(\text{messi}) &\rightsquigarrow e_4 \\ \text{playsfor}(\text{villa}, \text{FCbarca}) &\rightsquigarrow e_3 & \text{Striker}(\text{villa}) &\rightsquigarrow e'_4 \\ \text{playsfor}(\text{pedro}, \text{FCbarca}) &\rightsquigarrow e_3 & \text{Striker}(\text{pedro}) &\rightsquigarrow e''_4 \end{aligned}$$

with $p(e_3) = 0.5$ and $p(e_4) = p(e'_4) = p(e''_4) = 0.8$. Due to the last three CIs in \mathcal{T} , each of messi , villa , pedro is an answer to the CQ $\exists y. \text{playsfor}(x, y) \wedge \text{SoccerClub}(y)$ with probability 0.9.

Related Models in Probabilistic Databases. Nowadays, there is an abundance of probabilistic data models that provide compact representation of distributions over potentially large sets of *possible worlds*, see [15,30,2] and the references therein. Our pABoxes can be viewed as an open world version of the probabilistic data model studied by Dalvi and Suciu in [9]. It is as a less succinct version of *c-tables*, a traditional data model for probabilistic databases due to Imielinski and Lipski [18]. Since we are working with an open world semantics, pABoxes instead represent a distribution over *possible world descriptions*. Each such description may have any number of models. Note that our semantics is similar to the semantics of (“type 2”) probabilistic first-order and description logics [17,26].

Dealing with Inconsistencies. Of course, some of the ABoxes \mathcal{A}_E might be inconsistent w.r.t. the TBox \mathcal{T} used. In this case, it may be undesirable to let them contribute to the probabilities of answers. For example, if we use the pABox

$$\text{Striker}(\text{messi}) \rightsquigarrow e_1 \quad \text{Goalie}(\text{messi}) \rightsquigarrow e_2$$

with $p(e_1) = 0.8$ and $p(e_2) = 0.3$ and the TBox $\text{Goalie} \sqcap \text{Striker} \sqsubseteq \perp$, then messi is an answer to the query $\text{SoccerClub}(x)$ with probability 0.24 while one would probably expect it to be zero (which is the result when the empty TBox is used). We follow Antova, Koch, and Olteanu and advocate a pragmatic solution based on *rescaling* [2]. More specifically, we remove those ABoxes \mathcal{A}_E that are inconsistent w.r.t. \mathcal{T} and rescale the remaining set of ABoxes so that they sum up to probability one. In other words, we set

$$\hat{p}_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) = \frac{p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) - p(\mathcal{A}, \mathcal{T} \models \perp)}{1 - p(\mathcal{A}, \mathcal{T} \models \perp)}$$

where \perp is a Boolean query that is entailed exactly by those ABoxes \mathcal{A} that are inconsistent w.r.t. \mathcal{T} . The rescaled probability $\hat{p}_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q)$ can be computed in PTIME when this is the case both for $p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q)$ and $p(\mathcal{A}, \mathcal{T} \models \perp)$. Note that

rescaling results in some effects that might be unexpected such as reducing the probability of messi to be an answer to $\text{Striker}(x)$ from 0.8 to ≈ 0.74 when the above TBox is added.

In the remainder of the paper, for simplicity we will only admit TBoxes \mathcal{T} such that all ABoxes \mathcal{A} are consistent w.r.t. \mathcal{T} .

4 Query Rewriting

The aim of this section is to show that *FO-rewriting*, a prominent approach to traditional OBDA, are fruitful also in the case of computing probabilities of certain answers in *probabilistic* OBDA. In particular, we use it to lift the PTIME vs. #P dichotomy result on probabilistic databases recently obtained by Dalvi, Suciu, and Schnaitter [8] to probabilistic OBDA in DL-Lite.

4.1 Lifting FO-Rewritings to probabilistic OBDA

A first-order query $q_{\mathcal{T}}$ is an *FO-rewriting* of a CQ q and an FO-TBox \mathcal{T} (i.e., a first-order theory) if $\text{cert}_{\mathcal{T}}(q, \mathcal{A}) = \text{ans}(q_{\mathcal{T}}, \mathcal{I}_{\mathcal{A}})$ for every ABox \mathcal{A} , where $\mathcal{I}_{\mathcal{A}}$ denotes the ABox \mathcal{A} viewed as an interpretation. The query rewriting approach to traditional OBDA consists of computing, given a CQ q and a TBox \mathcal{T} , an FO-rewriting $q_{\mathcal{T}}$ and then handing it over for execution to a relational database system that stores the ABox \mathcal{A} .

The following observation states that FO-rewritings from traditional OBDA are also useful in probabilistic OBDA. We use $p_{\mathcal{A}}^d(\mathbf{a} \in q)$ to denote the probability that \mathbf{a} is an answer to the query q given the pABox \mathcal{A} viewed as a probabilistic database in the sense of Dalvi and Suciu [8]. More specifically,

$$p_{\mathcal{A}}^d(\mathbf{a} \in q) = \sum_{E \subseteq E_{\mathcal{A}} \mid \mathbf{a} \in \text{ans}(q, \mathcal{I}_{E})} p(E)$$

The following is immediate from the definitions.

Theorem 1 (Lifting). *Let \mathcal{T} be an FO-TBox, \mathcal{A} a pABox, q an n -ary CQ, $\mathbf{a} \in \text{Ind}(\mathcal{A})^n$ a candidate answer for q , and $q_{\mathcal{T}}$ an FO-rewriting of q relative to \mathcal{T} . Then $p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) = p_{\mathcal{A}}^d(\mathbf{a} \in q_{\mathcal{T}})$.*

From an application perspective, Theorem 1 enables the use of probabilistic database systems such as MayBMS, Trio, and MystiQ for implementing probabilistic OBDA [1,35,5]. Note that it might be necessary to adapt pABoxes in an appropriate way in order to match the data models of these systems. However, such modifications do not impair applicability of Theorem 1.

From a theoretical viewpoint, Theorem 1 establishes query rewriting as a useful tool for analyzing data complexity in probabilistic OBDA. We say that a CQ q is in PTIME relative to a TBox \mathcal{T} if there is a polytime algorithm that, given an ABox \mathcal{A} and a candidate answer $\mathbf{a} \in \text{Ind}(\mathcal{A})^n$ to q , computes $p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q)$. We say that q is #P-hard relative to \mathcal{T} if the afore mentioned

problem is hard for the counting complexity class #P [33]. We pursue a non-uniform approach to the complexity of query answering in probabilistic OBDA, as recently initiated in [27]: ideally, we would like to understand the precise complexity of every CQ q relative to every TBox \mathcal{T} , against the background of some preferably expressive ‘master logic’ used for \mathcal{T} .

Unsurprisingly, pABoxes are too strong a formalism to admit *any* tractable queries worth mentioning. An n -ary CQ q is *trivial* for a TBox \mathcal{T} iff for every ABox \mathcal{A} , we have $\text{cert}_{\mathcal{T}}(\mathcal{A}, q) = \text{Ind}(\mathcal{A})^n$.

Theorem 2. *Over pABoxes, every CQ q is #P-hard relative to every first-order TBox \mathcal{T} for which it is nontrivial.*

Theorem 2 motivates the study of more lightweight probabilistic ABox formalisms. While pABoxes (roughly) correspond to c-tables, which are among the most expressive probabilistic data models, we now move to the other end of the spectrum and introduce ipABoxes as a counterpart of *tuple independent databases* [9,12]. Arguably, the latter are the most inexpensive probabilistic data model that is still useful.

Definition 3 (ipABox). *An assertion-independent pABox (or ipABox) is a probabilistic ABox in which all event expressions are atomic and where each atomic event expression is associated with at most one ABox assertion.*

To save notation, we write ipABoxes in the form (\mathcal{A}, p) where \mathcal{A} is an ABox and p is a map $\mathcal{A} \rightarrow [0, 1]$ that assigns a probability to each ABox assertion. In this representation, the events are implicit (one atomic event per ABox assertion) and we write $p(\alpha)$ to denote the probability of the event associated with the assertion α . Analogously to pABoxes, all events (and thus all assertions) are independent. Note that the answer probability of $\mathbf{a} \in \text{Ind}(\mathcal{A})^n$ to an n -ary conjunctive query q over an ipABox (\mathcal{A}, p) relative to a TBox \mathcal{T} simplifies to

$$p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) = \sum_{\mathcal{A}' \subseteq \mathcal{A} : \mathbf{a} \in \text{cert}_{\mathcal{T}}(q, \mathcal{A}')} p(\mathcal{A}')$$

where $p(\mathcal{A}') = \prod_{\alpha \in \mathcal{A}'} p(\alpha) \cdot \prod_{\alpha \in \mathcal{A} \setminus \mathcal{A}'} (1 - p(\alpha))$.

Reconsidering our web data extraction example, it turns out that cases (1) and (4) yield ipABoxes, whereas cases (2), (3), and (5) do not. We refer to [32] for a discussion of the usefulness of ipABoxes/tuple independent databases. For the remainder of the paper, we assume that only ipABoxes are admitted unless explicitly noted otherwise.

4.2 Lifting the PTIME vs. #P Dichotomy

We now use Theorem 1 to lift a PTIME vs. #P dichotomy recently obtained in the area of probabilistic databases to probabilistic OBDA in DL-Lite. Note that, for any CQ and DL-Lite TBox, an FO-rewriting is guaranteed to exist [6]. The central observation is that, by Theorem 1, computing the probability of answers to a CQ q relative to a TBox \mathcal{T} over ipABoxes is *exactly the same problem* as computing the probability of answers to $q_{\mathcal{T}}$ over (ipABoxes viewed as) tuple




Fig. 1. Example queries

independent databases. We can thus analyze the complexity of CQs/TBoxes over ipABoxes by analyzing the complexity of their rewritings. In particular, standard rewriting techniques produce for each CQ and DL-Lite TBox an FO-rewriting that is a union of conjunctive queries (a UCQ) and thus, together with Theorem 1, Dalvi, Suciu and Schnaitter’s PTIME vs. #P dichotomy for UCQs over tuple independent databases [8] immediately yields the following.

Theorem 3 (Abstract Dichotomy). *Let q be a CQ and \mathcal{T} a DL-Lite TBox. Then q is in PTIME relative to \mathcal{T} or q is #P-hard relative to \mathcal{T} .*

Note that Theorem 3 actually holds for *every* DL that enjoys FO-rewritability. Although interesting from a theoretical perspective, Theorem 3 is not fully satisfactory as it does not tell us *which* CQs are in PTIME relative to which TBoxes. In the remainder of this section, we carry out a careful inspection of the FO-rewritings obtained in our framework and of the dichotomy result obtained by Dalvi, Suciu and Schnaitter, which results in a more concrete formulation of the dichotomy stated in Theorem 3 and provides a transparent characterization of the PTIME cases. For simplicity, we concentrate on CQs that are connected, Boolean, and do not contain individual names.

For two CQs q, q' and a TBox \mathcal{T} , we say that q \mathcal{T} -implies q' and write $q \sqsubseteq_{\mathcal{T}} q'$ when $\text{cert}_{\mathcal{T}}(q, \mathcal{A}) \subseteq \text{cert}_{\mathcal{T}}(q', \mathcal{A})$ for all ABoxes \mathcal{A} ; we say that q and q' are \mathcal{T} -equivalent and write $q \equiv_{\mathcal{T}} q'$ if $q \sqsubseteq_{\mathcal{T}} q'$ and $q' \sqsubseteq_{\mathcal{T}} q$; we say that q is \mathcal{T} -minimal if there is no $q' \subsetneq q$ such that $q \equiv_{\mathcal{T}} q'$. When \mathcal{T} is empty, we simply drop it from the introduced notation, writing for example $q \sqsubseteq q'$ and speaking of *minimality*. To have more control over the effect of the TBox, we will generally work with CQs q and TBoxes \mathcal{T} such that q is \mathcal{T} -minimal. This is without loss of generality because for every CQ q and TBox \mathcal{T} , we can find a CQ q' that is \mathcal{T} -minimal and such that $q \equiv_{\mathcal{T}} q'$ [4]; note that the answer probabilities relative to \mathcal{T} are identical for q and q' .

We now introduce a class of queries that will play a crucial role in our analysis.

Definition 4 (Simple Tree Queries). *A CQ q is a simple tree if there is a variable $x_r \in \text{var}(q)$ that occurs in every atom in q , i.e., all atoms in q are of the form $A(x_r)$, $r(x_r, y)$, or $r(y, x_r)$ ($y = x_r$ is possible). Such a variable x_r is called a root variable.*

As examples, consider the CQs in Figure 1, which are all simple tree queries. The following result shows why simple tree queries are important. A UCQ \hat{q} is *reduced* if for all disjuncts q, q' of \hat{q} , $q \sqsubseteq q'$ implies $q = q'$.

Theorem 4. *Let q be a CQ and \mathcal{T} a DL-Lite TBox such that q is \mathcal{T} -minimal and not a simple tree query. Then q is #P-hard relative to \mathcal{T} .*

Proof. (sketch) Let $q_{\mathcal{T}}$ be a UCQ that is an FO-rewriting of q relative to \mathcal{T} . By definition of FO-rewritings, we can w.l.o.g. assume that q occurs as a disjunct of $q_{\mathcal{T}}$. The following is shown in [8]:

1. if a minimal CQ does not contain a variable that occurs in all atoms, then it is #P-hard over tuple independent databases;
2. if a reduced UCQ \hat{q} contains a CQ that is #P-hard over tuple independent databases, then \hat{q} is also hard over tuple independent databases.

Note that since q is \mathcal{T} -minimal, it is also minimal. By Points 1 and 2 above, it thus suffices to show that $q_{\mathcal{T}}$ can be converted into an equivalent *reduced* UCQ such that q is still a disjunct, which amounts to proving that there is no disjunct q' in $q_{\mathcal{T}}$ such that $q \sqsubseteq q'$ and $q' \not\sqsubseteq q$. The details of the proof, which is surprisingly subtle, are given in the appendix. \square

To finish analyzing the dichotomy, it thus remains to analyze simple tree queries. We say that a role R is \mathcal{T} -generated in a CQ q if one of the following holds: (i) there is an atom $R(x_r, y) \in q$ and $y \neq x_r$; (ii) there is an atom $A(x_r) \in q$ and $\mathcal{T} \models \exists R \sqsubseteq A$; (iii) there is an atom $S(x, y) \in q$ with x a root variable and such that $y \neq x$ occurs only in this atom, and $\mathcal{T} \models \exists R \sqsubseteq \exists S$. The concrete version of our dichotomy result is as follows. Its proof is based on a careful analysis of FO-rewritings and the results in [10].

Theorem 5 (Concrete Dichotomy). *Let \mathcal{T} be a DL-Lite TBox. A \mathcal{T} -minimal CQ q is in PTIME relative to \mathcal{T} iff*

1. q is a simple tree query, and
2. if r and r^- are \mathcal{T} -generated in q , then $\{r(x, y)\} \sqsubseteq_{\mathcal{T}} q$ or q is of the form $\{S_1(x, y), \dots, S_k(x, y)\}$ for roles S_1, \dots, S_k .

Otherwise, q is #P-hard relative to \mathcal{T} .

As examples, consider again the queries q_1 , q_2 , and q_3 in Figure 1 and let \mathcal{T}_\emptyset be the empty TBox. All CQs are \mathcal{T}_\emptyset -minimal, q_1 and q_2 are in PTIME, and q_3 is #P-hard (all relative to \mathcal{T}_\emptyset). Now consider the TBox $\mathcal{T} = \{\exists s \sqsubseteq \exists r\}$. Then q_1 is \mathcal{T} -minimal and still in PTIME; q_2 is \mathcal{T} -minimal, and is now #P-hard because both s and s^- is \mathcal{T} -generated. The CQ q_3 can be made \mathcal{T} -minimal by dropping the r -atom, and is in PTIME relative to \mathcal{T} .

Approximations. Theorems 4 and 5 show that only very simple pairs (q, \mathcal{T}) can be answered in PTIME. Hence, it is reasonable to consider the approximation of answer probabilities. Of particular relevance are *fully-polynomial randomized approximation schemes (FPRASes)*, see [22,34] for a definition and more details. In [9] it is observed that there is an FPRAS for every UCQ over a probabilistic database. As a consequence of Theorem 1, there is thus also an FPRAS for every pair (q, \mathcal{T}) such that q is FO-rewritable relative to \mathcal{T} ; in particular, there are FPRASes for every CQ and DL-Lite ontology even if additional features such

as role inclusions are admitted. This observation clearly gives hope for practical feasibility of probabilistic OBDA. In the full paper, we also consider the existence of FPRASes for more expressive ontology languages [19].

5 Beyond Query Rewriting

A CQ q is *FO-rewritable* relative to a TBox \mathcal{T} if there is an FO-rewriting of q relative to \mathcal{T} . The aim of this section is to establish that, in a sense, FO-rewriting is a *complete* tool for proving PTIME results for CQ answering in probabilistic OBDA: we show that whenever a CQ q is not FO-rewritable relative to a TBox \mathcal{T} , then q is #P-hard relative to \mathcal{T} ; thus, when a query is in PTIME relative to a TBox \mathcal{T} , then this can *always* be shown via FO-rewritability. To achieve this goal, we select \mathcal{ELI} as the TBox language because it properly generalizes DL-Lite (as in the previous sections we disregard \perp) and, unlike DL-Lite, also embraces non FO-rewritable CQs/TBoxes. Note that, in traditional OBDA, there is a drastic difference in data complexity of CQ-answering between DL-Lite and \mathcal{ELI} : the former is in AC_0 while the latter is PTIME-complete.

We focus on Boolean CQs q that are *rooted*, i.e., q involves at least one individual name and is connected. This is a natural case since, for any non-Boolean connected CQ $q(\mathbf{x})$ and potential answer \mathbf{a} , the probability $p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q(\mathbf{x}))$ that \mathbf{a} is a certain answer to q w.r.t. \mathcal{A} and \mathcal{T} is identical to the probability $p(\mathcal{A}, \mathcal{T} \models q[\mathbf{a}])$ that \mathcal{A} and \mathcal{T} entail the rooted Boolean CQ $q[\mathbf{a}]$. The following theorem says that there is no hope for PTIME algorithms in the case a query q is not FO-rewritable relative to \mathcal{T} .

Theorem 6. *If a Boolean rooted CQ q is not FO-rewritable relative to an \mathcal{ELI} -TBox \mathcal{T} , then q is #P-hard relative to \mathcal{T} .*

As a by-product of Theorem 6 we obtain the following dichotomy.

Theorem 7 (\mathcal{ELI} dichotomy). *Let q be a rooted Boolean CQ and \mathcal{T} an \mathcal{ELI} -TBox. Then q is in PTIME relative to \mathcal{T} or #P-hard relative to \mathcal{T} .*

6 Conclusion

We have introduced a framework for ontology-based access to probabilistic data that can be implemented using existing probabilistic database system, and we have analyzed the data complexity of computing answer probabilities in this framework. There are various opportunities for future work. For example, it would be interesting to extend the *concrete* dichotomy: on the one hand, it can be extended to CQs that involve constants and are not necessarily connected; on the other hand, one could study more expressive versions of DL-Lite that, for example, allow role hierarchy statements in the TBox. It would also be worthwhile to add means to express uncertainty to the TBox formalism instead of admitting it only in the ABox; this is done for example in [28,12], but it remains to be seen whether the semantics used there is appropriate for our purposes.

References

1. Antova, L., Jansen, T., Koch, C., Olteanu, D.: Fast and simple relational processing of uncertain data. In: Proc. of ICDE. 983–992 (2008)
2. Antova, L., Koch, C., Olteanu, D.: 10^{10^6} worlds and beyond: efficient representation and processing of incomplete information. VLDB J. 18(5), 1021–1040 (2009)
3. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook. Cambridge University Press (2003)
4. Bienvenu, M., Lutz, C., Wolter, F.: Query containment in description logics reconsidered. In: Proc. of KR (2012)
5. Boulos, J., Dalvi, N.N., Mandhani, B., Mathur, S., Ré, C., Suciu, D.: MYSTIQ: a system for finding more answers by using probabilities. In: Proc. of SIGMOD. 891–893 (2005)
6. Calvanese, D., Giacomo, G.D., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. J. Autom. Reasoning 39(3), 385–429 (2007)
7. Dalvi, N.N., Ré, C., Suciu, D.: Probabilistic databases: diamonds in the dirt. Commun. ACM 52(7), 86–94 (2009)
8. Dalvi, N.N., Schnaitter, K., Suciu, D.: Computing query probability with incidence algebras. In: Proc. of PODS. 203–214. ACM (2010)
9. Dalvi, N.N., Suciu, D.: Efficient query evaluation on probabilistic databases. VLDB J. 16(4), 523–544 (2007)
10. Dalvi, N.N., Suciu, D.: The dichotomy of probabilistic inference for unions of conjunctive queries. J. ACM 59(6), 30 (2012)
11. Finger, M., Wassermann, R., Cozman, F.G.: Satisfiability in \mathcal{EL} with sets of probabilistic ABoxes. Proc. of DL. CEUR-WS, Vol. 745 (2011)
12. Fuhr, N., Rölleke, T.: A probabilistic relational algebra for the integration of information retrieval and database systems. ACM Trans. Inf. Syst. 15(1), 32–66 (1997)
13. Furche, T., Gottlob, G., Grasso, G., Gunes, O., Guo, X., Kravchenko, A., Orsi, G., Schallhart, C., Sellers, A.J., Wang, C.: Diadem: domain-centric, intelligent, automated data extraction methodology. In Proc. of WWW. 267–270. ACM (2012)
14. Gottlob, G., Lukasiewicz, T., Simari, G.I.: CQ answering in probabilistic datalog+/- ontologies. In Proc. of RR. Vol. 6902 of LNCS, 77–92. Springer (2011)
15. Green, T.J., Tannen, V.: Models for incomplete and probabilistic information. IEEE Data Engineering Bulletin 29(1), 17–24 (2006)
16. Gupta, R., Sarawagi, S.: Creating probabilistic databases from information extraction models. In Proc. of VLDB. 965–976. ACM (2006)
17. Halpern, J.Y.: An analysis of first-order logics of probability. Artif. Intell. 46(3), 311–350 (1990)
18. Imielinski, T., Jr., W.L.: Incomplete information in relational databases. J. of the ACM 31(4), 761–791 (1984)
19. Jung, J.C., Lutz, C.: Ontology-based access to probabilistic data with OWL QL. In: Proc. of ISWC. pp. 182–197. Springer (2012)
20. Jerrum, M., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. Theor. Comput. Sci. 43, 169–188 (1986)
21. Karger, D.R.: A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem. SIAM J. Comput. 29(2), 492–514 (1999)
22. Karp, R.M., Luby, M.: Monte-carlo algorithms for enumeration and reliability problems. In Proc. of FoCS. 56–64. IEEE Computer Society (1983)

23. Kontchakov, R., Lutz, C., Toman, D., Wolter, F., Zakharyashev, M.: The combined approach to query answering in DL-Lite. In Proc. of KR. AAAI Press (2010)
24. Laender, A.H.F., Ribeiro-Neto, B.A., da Silva, A.S., Teixeira, J.S.: A brief survey of web data extraction tools. SIGMOD Record 31(2), 84–93 (2002)
25. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the semantic web. J. Web Sem. 6(4), 291–308 (2008)
26. Lutz, C., Schröder, L.: Probabilistic description logics for subjective uncertainty. In Proc. of KR. AAAI Press (2010)
27. Lutz, C., Wolter, F.: Non-uniform data complexity of query answering in description logics. In: Proc. of KR. AAAI Press (2012)
28. Raedt, L.D., Kimmig, A., Toivonen, H.: Problog: a probabilistic prolog and its application in link discovery. In Proc. of IJCAI. 2468–2473. AAAI Press (2007)
29. Rossman, B.: Homomorphism preservation theorems. J. ACM. 55(3). 1–54 (2008).
30. Sarma, A.D., Benjelloun, O., Halevy, A.Y., Widom, J.: Working models for uncertain data. In: Proc. of ICDE. IEEE Computer Society (2006)
31. Straccia, U.: Top-k retrieval for ontology mediated access to relational databases. In: Information Sciences. 108, 1–23 (2012).
32. Suciú, D., Olteanu, D., Ré, C., Koch, C.: Probabilistic Databases. Synthesis Lectures on Data Management, Morgan & Claypool Publishers (2011)
33. Valiant, L.G.: The complexity of enumeration and reliability problems. SIAM J. Comput. 8(3), 410–421 (1979)
34. Vazirani, V.V.: Approximation algorithms. Springer (2001)
35. Widom, J.: Trio: A system for integrated management of data, accuracy, and lineage. In Proc. of CIDR. 262–276 (2005)
36. Zenklusen, R., Laumanns, M.: High-confidence estimation of small s - t reliabilities in directed acyclic networks. Networks 57(4), 376–388 (2011)

A Proofs for Section 4

A.1 Theorem 2

Theorem 2. *Over pABoxes, every CQ q is #P-hard relative to every first-order TBox \mathcal{T} for which it is nontrivial.*

Proof. The proof is by reduction of counting the number of satisfying assignments of a propositional formula.¹ Assume that q has answer variables x_1, \dots, x_n and let φ be a propositional formula over variables z_1, \dots, z_m . Convert φ into a pABox \mathcal{A} as follows: take q viewed as an ABox, replacing every variable x with an individual name a_x ; then associate every ABox assertion with φ viewed as an event expression over events z_1, \dots, z_m and set $p(z_i) = 0.5$ for all i . We are interested in the answer $\mathbf{a} = a_{x_1} \cdots a_{x_n}$. For all $E \subseteq E_{\mathcal{A}}$ with $E \not\models \varphi$, we have $\mathcal{A}_E = \emptyset$ and thus $\mathbf{a} \notin \text{cert}_{\mathcal{T}}(q, \mathcal{A}_E)$ since q is non-trivial for \mathcal{T} . For all $E \subseteq E_{\mathcal{A}}$ with $E \models \varphi$, the ABox \mathcal{A}_E is the ABox-representation of q and thus $\mathbf{a} \in \text{cert}_{\mathcal{T}}(q, \mathcal{A}_E)$. Consequently, the number of assignments that satisfy φ is $p_{\mathcal{A}, \mathcal{T}}(\mathbf{a} \in q) \cdot 2^m$. Thus, there is a PTIME algorithm for counting the number of satisfying assignments given an oracle for computing answer probabilities for q and \mathcal{T} . \square

A.2 Theorem 4

Throughout the whole section, we assume CQs to be Boolean and connected, and to not contain individual names. In this section and the subsequent ones, we will sometimes make use of canonical models as defined in [23], introduced here as a preliminary. To construct the *canonical model* $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ of an ABox \mathcal{A} and a DL-Lite TBox \mathcal{T} , we start with \mathcal{A} viewed as an interpretation and then exhaustively apply the CIs from \mathcal{T} as rules, introducing fresh elements for existential quantifiers. Formally, the domain of $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ consists of paths of the form $aR_1 \cdots R_n$, $n \geq 0$, such that the following conditions hold:

(agen) $\mathcal{A}, \mathcal{T} \models \exists R_1(a)$ but $R_1(a, b) \notin \mathcal{A}$ for all $b \in \text{Ind}(\mathcal{A})$

(rgen) for $i < n$, $\mathcal{T} \models \exists R_i^- \sqsubseteq \exists R_{i+1}$ and $R_i^- \neq R_{i+1}$ (written $c_{R_i} \rightsquigarrow c_{R_{i+1}}$).

We denote by $\text{tail}(\sigma)$ the last element in a path σ . Now, $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ is defined as follows:

$$\begin{aligned} \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \{a \cdot c_{R_1} \cdots c_{R_n} \mid a \in \text{Ind}(\mathcal{A}), n \geq 0, a \rightsquigarrow c_{R_1} \rightsquigarrow \cdots \rightsquigarrow c_{R_n}\}, \\ a^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= a, \text{ for all } a \in \text{Ind}(\mathcal{A}), \\ A^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \{a \in \text{Ind}(\mathcal{A}) \mid \mathcal{K} \models A(a)\} \cup \{\sigma \cdot R \in \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} \mid \mathcal{T} \models \exists R^- \sqsubseteq A\}, \\ P^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \{(a, b) \in \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \mid P(a, b) \in \mathcal{A}\} \cup \\ &\quad \{(\sigma, \sigma \cdot c_P) \in \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} \times \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} \mid \text{tail}(\sigma) \rightsquigarrow c_P\} \cup \\ &\quad \{(\sigma \cdot c_{P^-}, \sigma) \in \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} \times \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} \mid \text{tail}(\sigma) \rightsquigarrow c_{P^-}\}, \end{aligned}$$

¹ Throughout the paper, we use the standard oracle-based notion of reduction originally introduced by Valiant in the context of counting complexity [33].

where ‘.’ denotes concatenation. Domain elements in $\Delta^{\mathcal{I}_{\mathcal{A},\mathcal{T}}}$ that are true path (i.e., do not only consists of an individual name) are called *anonymous*. The following is the central property of canonical models.

Theorem 8 ([23]). *For every consistent DL-Lite KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and every CQ q , we have $\mathbf{a} \in \text{cert}_{\mathcal{T}}(q, \mathcal{A})$ iff $\mathcal{I}_{\mathcal{A},\mathcal{T}} \models q[\mathbf{a}]$.*

We will sometimes also use canonical models $\mathcal{I}_{q,\mathcal{T}}$ for a CQ q and a TBox \mathcal{T} , defined as $\mathcal{I}_{\mathcal{A}_q,\mathcal{T}}$ where \mathcal{A}_q is q viewed as an ABox, i.e., the variables in q are viewed as the ABox individuals of \mathcal{A}_q . In particular, we thus have

$$\Delta^{\mathcal{I}_{q,\mathcal{T}}} = \{x \cdot c_{R_1} \cdots c_{R_n} \mid x \in \text{var}(q), n \geq 0, x \rightsquigarrow c_{R_1} \rightsquigarrow \cdots \rightsquigarrow c_{R_n} \text{ in } \mathcal{A}_q\}.$$

The following is easy to prove.

Lemma 1 ([4]). *$q \sqsubseteq_{\mathcal{T}} q'$ iff there is a homomorphism from q' to $\mathcal{I}_{q,\mathcal{T}}$, i.e., a map $h : \text{var}(q') \rightarrow \Delta^{\mathcal{I}_{q,\mathcal{T}}}$ such that the following conditions are satisfied:*

1. *if $A(t) \in q'$, then $h(t) \in A^{\mathcal{I}_{q,\mathcal{T}}}$;*
2. *if $r(t, t') \in q'$, then $(h(t), h(t')) \in r^{\mathcal{I}_{q,\mathcal{T}}}$.*

As explained in the proof sketch in the main paper, Theorem 4 is a consequence of the following lemma. We call an FO-rewriting of a CQ q and a TBox \mathcal{T} that happens to be a UCQ a *UCQ-rewriting*.

Lemma 2. *Let q be a CQ and \mathcal{T} a DL-Lite TBox such that q is \mathcal{T} -minimal. Then there is a UCQ-rewriting $q_{\mathcal{T}}$ of q relative to \mathcal{T} that is reduced and such that q occurs as a disjunct in $q_{\mathcal{T}}$. Moreover, if q is a simple tree query, then so is every disjunct of $q_{\mathcal{T}}$.*

Proof. Let $q_{\mathcal{T}}$ be a UCQ-rewriting of q relative to \mathcal{T} . It is easily verified that standard approaches such as the one from [6] will generate a UCQ in which every disjunct is a simple tree query if q is a simple tree query. By definition of FO-rewritings, $q_{\mathcal{T}} \vee q$ is also an FO-rewriting and thus we can w.l.o.g. assume that $q_{\mathcal{T}}$ contains q as a disjunct. Since q is \mathcal{T} -minimal and thus minimal, we can further assume that every disjunct of $q_{\mathcal{T}}$ is minimal. We show that q is not eliminated when $q_{\mathcal{T}}$ is converted into reduced form, i.e., that for every disjunct q' of $q_{\mathcal{T}}$ with $q \sqsubseteq q'$, we have $q \equiv q'$. Choose such a disjunct q' . Using the definition of FO-rewritings, it is not hard to show that we must have $q' \sqsubseteq_{\mathcal{T}} q$.

We first consider the case where q has two root variables. Then q must be of the form $\{s_1(x, y), \dots, s_k(x, y)\}$, where the s_i are role names or inverse roles. If $k > 1$, then $q' \sqsubseteq_{\mathcal{T}} q$ implies that $q' = q$ or $q' = \{s_1(x, x), \dots, s_k(x, x)\}$. In both cases, we obtain $q' \sqsubseteq q$, which yields $q \equiv q'$. If $k = 1$, then q has the form $s(x, y)$ with $x \neq y$ and since $q \sqsubseteq q'$ and the CQ q' (like any CQ) cannot be empty, we have $q' = q$ and are done.

Now assume that q has only a single root variable. As a preliminary, we make the following observation.

Claim. If x' is a root variable of q' , then there is a homomorphism from q' to q

that maps x' to the unique root variable of q .

Proof of claim. Since $q \sqsubseteq q'$, there is a homomorphism g from q' to q . Let x' be a root variable in q' and x the root of q . We want to show that $g(x') = x$. Assume to the contrary that this is not the case. Then $g(x') = y \neq x$. Since q has only one root variable, y is not a root variable and q does neither contain concept atoms nor reflexive role atoms on y . Since g is a homomorphism, q' thus does not contain concept atoms or reflexive role atoms on x' . Since q and q' are simple tree queries, all variables from $\text{var}(q') \setminus \{x'\}$ are connected via some role atom to x' in q' and the only variable in q connected via some role atom to y is the root x . Therefore, we have $g(y') = x$ for all $y' \in \text{var}(q') \setminus \{x'\}$.

- There is a role R such that for all atoms $S(y', x')$ in q' , we have $R = S$.
Since q' is minimal, it then has the form $R(y', x')$. Due to the existence of g , there is an atom $R(x, y)$ in q . Since $q' \sqsubseteq_{\mathcal{T}} q$, we have $R(x, y) \sqsubseteq_{\mathcal{T}} q$. Since q is \mathcal{T} -minimal, this yields $q = R(x, y)$. This is a contradiction to q having only one root variable.
- There are atoms $R_1(y'_1, x')$ and $R_2(y'_2, x')$ in q' such that $R_1 \neq R_2$.
For all atoms $R(y', x')$ in q' , we have $R(x, y) \in q$ since $g(x') = y$ and $g(y') = x$. Since there is more than one role name R that occurs in such an atom, the restriction of q to the query q_y containing those atoms from q that involve y is not tree-shaped. Since $q' \sqsubseteq_{\mathcal{T}} q$, there is a homomorphism h from q to $\mathcal{I}_{q', \mathcal{T}}$. By construction of the latter and since q_y is not tree-shaped, the restriction of h to q_y is a homomorphism from q_y to q' . Thus, there is a y'_0 in q' such that for all atoms $R(y', x') \in q'$, we have $R(y'_0, x') \in q'$. Since q' is minimal, it thus cannot contain any other variables than y'_0 and x' .
Since y is not a root variable of q , there must be at least one atom in q that involves x , but not y . Let q^- be obtained from q by dropping all such atoms. Then g is also a homomorphism from q' to q^- . By construction of $\mathcal{I}_{q^-, \mathcal{T}}$, we find a homomorphism g^+ from $\mathcal{I}_{q', \mathcal{T}}$ to $\mathcal{I}_{q^-, \mathcal{T}}$ with $g^+(x') = y$ and $g^+(y'_0) = x$. Composing the homomorphism h , restricted to q^- , with g^+ yields a homomorphism from q to $\mathcal{I}_{q^-, \mathcal{T}}$. Consequently, $q' \sqsubseteq_{\mathcal{T}} q$ in contradiction to the minimality of q .

This finishes the proof of the claim.

Since $q' \sqsubseteq_{\mathcal{T}} q$, there is a homomorphism h from q to $\mathcal{I}_{q', \mathcal{T}}$. Define the subset $Z \subseteq q'$ of atoms that are ‘hit by h ’ as

$$Z = \{A(x') \in q' \mid \exists A(x) \in q : h(x) = x'\} \cup \{r(x', y') \in q' \mid \exists r(x, y) \in q : h(x) = x' \wedge h(y) = y'\}.$$

We distinguish the following cases.

- $A(x) \in q$ implies $A(h(x)) \in Z$ and $r(x, y) \in q$ implies $r(h(x), h(y)) \in Z$.
Then $q' \sqsubseteq q$ and thus $q \equiv q'$. Consequently, we are done.
- There is an $A(x) \in q$ with $A(h(x)) \notin Z$.
Then x is the root variable in the simple tree query q . Since $h(x) \in A^{\mathcal{I}_{q', \mathcal{T}}}$, one of the following cases applies:

1. $h(x)$ is anonymous in $\mathcal{I}_{q,\mathcal{T}}$.
 Let $h(x) = x'c_{R_1} \cdots c_{R_k}$ with $k > 0$. Since q is a simple tree query, it follows that h maps all variables in q to the subtree of $\mathcal{I}_{q',\mathcal{T}}$ rooted at x' , i.e., all elements in the range of h have x' as a (not necessarily proper) prefix. Moreover, since the root variable x is not mapped to x' and all concept atoms in q involve x , there is even a match of q in the subtree rooted in $\mathcal{I}_{q',\mathcal{T}}$ at x' when all concept atoms at the root of that subtree are removed. By construction of $\mathcal{I}_{q',\mathcal{T}}$, there is thus a single atom at in q' such that q has a match in $\mathcal{I}_{\text{at},\mathcal{T}}$. Since $q \sqsubseteq q'$, we must find at also in q , modulo renaming of variables. Since q is \mathcal{T} -minimal, this implies $q = \{\text{at}\}$. Since $q \sqsubseteq q'$ and q' is non-empty and minimal, we must have $q' = \text{at}$, thus $q = q'$ and we are done.
 2. There is a $B(h(x)) \in q'$ such that $\mathcal{T} \models B \sqsubseteq A$.
 Then $h(x)$ is a root variable of q' . By the above claim, there is a homomorphism g from q' to q with $g(h(x)) = x$ and thus $B(x) \in q$. Since $\mathcal{T} \models B \sqsubseteq A$, this is a contradiction against \mathcal{T} -minimality of q and the fact that $A(x) \in q$.
 3. There is an $R(h(x), y') \in q'$ such that $\mathcal{T} \models \exists R \sqsubseteq A$.
 If $h(x)$ is a root variable of q' , we can argue as in the previous case. Otherwise, h must map all variables in q to $h(x)$, a root variable y' of q' with $S(y', h(x)) \in q'$ for at least one S , and to the anonymous subtree rooted at $h(x)$. Let $X' = \{R_1(y', h(x)), \dots, R_m(y', h(x))\}$ be all assertions in q' of this form. By construction of $\mathcal{I}_{q',\mathcal{T}}$, there is a match of q in $\mathcal{I}_{X',\mathcal{T}}$ and thus $X', \mathcal{T} \models q$. Since $q \sqsubseteq q'$, we find X' as a subquery X in q , probably after identifying the two variables in X . Since $X, \mathcal{T} \models q$, this is a contradiction against \mathcal{T} -minimality of q and the fact that $A(x) \in q \setminus X$.
- There is an $R(x, y) \in q$ with $R(h(x), h(y)) \notin Z$.
 Then either x or y is a root variable. We assume w.l.o.g. that x is the root variable, the other case is symmetric. Make a case distinction as follows:
1. $h(x)$ is anonymous in $\mathcal{I}_{q,\mathcal{T}}$.
 Then we can argue as in Case 1 above.
 2. $h(x)$ is a root variable of q' .
 Since $(h(x), h(y)) \in R^{\mathcal{I}_{q,\mathcal{T}}} \setminus Z$ and $h(x)$ is not anonymous, $h(y)$ must be anonymous by construction of $\mathcal{I}_{q',\mathcal{T}}$. Thus, $(h(x), h(y)) \in S^{\mathcal{I}_{q',\mathcal{T}}}$ implies $S = R$, which yields that $S(x, y) \in q$ implies $S = R$. Moreover, q does not contain any concept atoms on y since y is not a root variable and thus, in summary, y is not involved in any other atoms than $R(x, y)$ in q . By the claim, there is a homomorphism g from q' to q with $g(h(x)) = x$. By construction of $\mathcal{I}_{q',\mathcal{T}}$, $(h(x), h(y)) \in R^{\mathcal{I}_{q',\mathcal{T}}}$ with $h(y)$ anonymous implies that there is no atom $R(h(x), z)$ in q' . Thus, g is still a homomorphism from q' to $q^- := q \setminus \{R(x, y)\}$. Consequently, there is an outgoing R -edge from x in $\mathcal{I}_{q^-, \mathcal{T}}$. Thus and since y is not involved in any other atoms than $R(x, y)$ in q , we have $q^- \sqsubseteq_{\mathcal{T}} q$, in contradiction to the \mathcal{T} -minimality of q .
 3. $h(x)$ is a non-root variable of q' .

Then, h must map all variables in q to $h(x)$, a root variable y' of q' with $S(y', h(x)) \in q'$ for at least one S , and to the anonymous subtree rooted at $h(x)$. Let $X' = \{R_1(y', h(x)), \dots, R_m(y', h(x))\}$ be all assertions in q' of this form. By Since $(h(x), h(y)) \in R^{\mathcal{I}_{q', \mathcal{T}}} \setminus Z$ and $h(x)$ is not anonymous, $h(y)$ must be anonymous. By construction of $\mathcal{I}_{q', \mathcal{T}}$, R^- is thus not among the R_1, \dots, R_m . Moreover, there is a match of q in $\mathcal{I}_{X', \mathcal{T}}$ and thus $X', \mathcal{T} \models q$. We can argue as in the previous case that y is not involved in any other atoms than $R(x, y)$ in q . Since $q \sqsubseteq q'$, we find X' as a subquery X of q , probably after identifying the two variables in X' . Clearly $X, \mathcal{T} \models q$. If y does not occur in X , then $X \subsetneq q$, in contradiction to the \mathcal{T} -minimality of q . Otherwise, since R^- is not among the R_1, \dots, R_m and y occurs in no other atoms than $R(x, y)$ in q , we must have $X' = R(y', h(x))$. From $X', \mathcal{T} \models q$, we thus obtain $\{R(x, y)\}, \mathcal{T} \models q$. Since q is \mathcal{T} -minimal and $R(x, y) \in q$, we must have $q = R(x, y)$. This is a contradiction to q having only one root variable. \square

A.3 Theorem 5

To prove Theorem 5, we analyze the results in [8,10]. As a preparation for that proof, we first introduce some notions and results, all taken from [10], to which we also refer for further explanations and examples.

A *disjunctive sentence* is a disjunction of *connected* CQs. A disjunctive sentence $q = q_1 \vee \dots \vee q_k$ is *symbol-connected* if the graph (V, E) with

- $V = \{q_1, \dots, q_k\}$
- $(q_i, q_j) \in E$ iff there is a relational symbol that occurs both in q_i and q_j

has precisely one connected component. A *root variable* of a UCQ q is a variable that occurs in all atoms of q . A root variable x of q is a *separator variable* of q if additionally for every relation name r that occurs in the query, there is a number i_r such that every atom in q with symbol r contains exactly one occurrence of x and that is on position i_r .

Let x_1, \dots, x_n be a set of variables. A set Θ of predicates of the form $x_i < x_j$ is called *consistent* if there is a total order $<^*$ such that $x_i < x_j \in \Theta$ implies $x_i <^* x_j$. A CQ q is *ranked* if the set

$$\{x < y \mid r(x, y) \in q\}$$

is consistent.

In what follows, we assume the disjunctive sentence q to be reduced and all disjuncts p of q to be minimal, i.e., there is no $p' \subsetneq p$ with $p' \sqsubseteq p$. We call a disjunctive sentence q *immediately unsafe* if it is symbol-connected, contains at least one variable, and does not have a separator.

Theorem 9 (Theorem 4.4 of [10]). *If a disjunctive query q is ranked and immediately unsafe, then computing $P(q)$ is #P-hard.*

The following also follows from this theorem and the algorithm in [10].

Corollary 1. *If a disjunctive query q is ranked and does not have a separator, then computing $P(q)$ is #P-hard.*

Proof. We can write q as $q_1 \vee \dots \vee q_k$ where q_1, \dots, q_k are symbol-connected disjunctive sentences. The algorithm in [10] computes $P(q)$ as $1 - (1 - P(q_1)) \cdot \dots \cdot (1 - P(q_k))$. Since q is not ranked, there exists i such that q_i does not have a separator. Hence, the algorithm fails in the computation of $P(q_i)$ because of Theorem 9. \square

In [10], the following is proven (note that we use slightly different notation, e.g., D denotes a tuple-independent probabilistic database).

Proposition 1. *Every UCQ q is computationally equivalent to a ranked UCQ \bar{q} over an extended vocabulary, i.e., the problem “given a probabilistic database D , compute $P_D(q)$ ” can be reduced in polynomial time to “given \bar{D} compute $P_{\bar{D}}(\bar{q})$ ”, and vice versa.*

Since the construction of \bar{q} in the proof of this Proposition is needed later, we present it here for the binary case. We define three functions $\tau_=, \tau_<, \tau_>$ as follows:

$$\begin{aligned}\tau_=(x_1, x_2) &= (x_1) \\ \tau_<(x_1, x_2) &= (x_1, x_2) \\ \tau_>(x_1, x_2) &= (x_2, x_1)\end{aligned}$$

Additionally, introduce for every binary relation symbol R occurring in q three fresh relation symbols $R^=$ (which is unary), $R^<$, and $R^>$ (both binary). Finally define $\Omega = \{=, <, >\}$.

If $p = p_0, R_1(y_1, z_1), \dots, R_m(y_m, z_m)$ is a CQ in q such that p_0 contains precisely the unary atoms of p and $\sim_1, \dots, \sim_m \in \Omega$ we call the CQ

$$p_0, \bigwedge_{i=1}^m R_i^{\sim_i}(\tau_{\sim_i}(y_i, z_i)), y_i \sim_i z_i \quad (1)$$

the \sim_1, \dots, \sim_m -reduct of p . We call the reduct *strict* if $\sim_i \in \{<, >\}$ for all $i \in \{1, \dots, m\}$.

The ranking \bar{q} of q is now the disjunction of all \sim_1, \dots, \sim_m -reducts of some disjunct p in q and $\sim_1, \dots, \sim_m \in \Omega$. Obviously, the set of appended order predicates $y_i \sim_i z_i$ might be inconsistent and therefore the corresponding reduct can be dropped from the ranking. However, it is easy to see that there is a consistent reduct of every disjunct p of a UCQ q . Moreover, there is even a consistent strict reduct of every p in q when one \sim_i is fixed (to $<$ or $>$). We prove that strict reducts are not removed during the reduction of the obtained UCQ.

Lemma 3. *If q is a reduced UCQ and p' a strict \sim_1, \dots, \sim_m -reduct of some CQ p in q . Then p' is contained in the ranking \bar{q} of q , i.e., p' is not removed during the reduction \bar{q} .*

Proof. Suppose to the contrary that there is some CQ p_1 in q and $\sim'_1, \dots, \sim'_k \in \Omega$ such that the \sim'_1, \dots, \sim'_k -reduct p'_1 of p_1 satisfies $p'_1 \sqsubseteq p'$ and $p'_1 \neq p'$. Thus, there is a homomorphism h from p' to p'_1 . We claim that h is also a homomorphism from p to p_1 . Let $p = p_0, R_1(y_1, z_1), \dots, R_m(y_m, z_m)$ such that p_0 contains precisely the unary atoms of p . Then we have:

- If $A(x) \in p$, formula (1) implies that $A(x) \in p'$ and since h is homomorphism, we have also $A(h(x)) \in p'_1$. Using again formula (1) yields $A(h(x)) \in p_1$.
- Let $R_i(x, y) \in p$. Since p' is a strict reduct, we have to distinguish two cases:
 - Let $\sim_i = <$. Then $R_i^<(x, y) \in p'$ and since h is a homomorphism, we have that $R_i^<(h(x), h(y)) \in p'_1$. By formula (1), $R_i(h(x), h(y)) \in p_1$.
 - Let $\sim_i = >$. Then $R_i^>(y, x) \in p'$ and since h is a homomorphism, we have that $R_i^>(h(y), h(x)) \in p'_1$. By formula (1), $R_i(h(x), h(y)) \in p_1$.

Because of the homomorphism from p to p_1 , we obtain $p_1 \sqsubseteq p$. Since q is reduced this yields $p_1 = p$ and minimality of p implies that the homomorphism h is surjective (from p to p_1). Hence, h is also a surjective homomorphism from p' to p'_1 . \square

With the next lemma we provide an intuition why the notion of “being \mathcal{T} -generated” is central in the presence of rewritings.

Lemma 4. *Let q be a simple tree query and \mathcal{T} a DL-Lite TBox such that q is \mathcal{T} -minimal. Then there exists a reduced UCQ-rewriting $q_{\mathcal{T}}$ of q relative to \mathcal{T} such that the following conditions are satisfied:*

1. every disjunct of $q_{\mathcal{T}}$ is a simple tree query;
2. q is a disjunct in $q_{\mathcal{T}}$;
3. for every role R the following are equivalent
 - R is \mathcal{T} -generated in q or $\{R(x, y)\} \sqsubseteq_{\mathcal{T}} q$
 - there exists a disjunct q' in $q_{\mathcal{T}}$ such that q' contains an atom $R(x, y)$ where x is a root variable of q' and $x \neq y$

Proof. Let q be a simple tree query and \mathcal{T} a DL-Lite TBox. The reduced UCQ-rewriting $q_{\mathcal{T}}$ whose existence is guaranteed by Lemma 2 satisfies Conditions 1 and 2 of Lemma 4. It thus remains to show that it also satisfies Condition 3. We deal with the two directions separately.

“ \Rightarrow ”: If $\{R(x, y)\} \sqsubseteq_{\mathcal{T}} q$, we can set $q'_{\mathcal{T}} \equiv q_{\mathcal{T}} \vee \{R(x, y)\}$ and it is clear that the disjunct $\{R(x, y)\}$ is not removed during reduction of $q'_{\mathcal{T}}$. If R is a role that is \mathcal{T} -generated in q , then one of the following cases applies:

- $R(x, y) \in q$, x is a root variable, and $x \neq y$.
By choice of $q_{\mathcal{T}}$, q is a disjunct of $q_{\mathcal{T}}$ and we are done.
- $A(x) \in q$ and $\mathcal{T} \models \exists R \sqsubseteq A$.
Let $q_R = (q \setminus \{A(x)\}) \cup \{R(x, y)\}$ with y a fresh variable. Clearly, we have $\mathcal{A}_{q_R}, \mathcal{T} \models q$, thus $\mathcal{I}_{\mathcal{A}_{q_R}} \models q_{\mathcal{T}}$. Let q' be a disjunct of $q_{\mathcal{T}}$ with $\mathcal{I}_{\mathcal{A}_{q_R}} \models q'$. Then $q_R \sqsubseteq q'$ and we find a homomorphism h from q' to q_R .
We show that there must be an atom $R(x', y') \in q'$ with $h(x') = x, h(y') = y$. Assume to the contrary that this is not the case. Let $q^- = q_R \setminus \{R(x, y)\}$

and observe that, since $A(x) \notin q^-$, we have $q^- \subsetneq q$. Then h is also a homomorphism from q' to q^- and thus $q^- \sqsubseteq q'$ and also $q \sqsubseteq q'$. Since $q_{\mathcal{T}}$ is reduced and both q and q' are disjuncts of $q_{\mathcal{T}}$, it follows from the latter that $q = q'$. Thus $q^- \sqsubseteq q'$ means in fact $q^- \sqsubseteq q$, which is a contradiction to the minimality of q .

Consider the atom $R(x', y') \in q'$ with $h(x') = x$, $h(y') = y$ whose existence we have just proved. Clearly, $x' \neq y'$. Thus, it remains to show that x' is a root variable in q' . Assume to the contrary that x' is not a root variable. Then y' must be a root variable and there must be an atom in q' that involves y' , but not x' . This implies that q contains an atom that involves y , but not x , in contradiction to x being a root variable of q .

- $S(x, y) \in q$ such that y occurs only once and $\mathcal{T} \models \exists S \sqsubseteq \exists R$.

Identical to previous case.

“ \Leftarrow ”: Assume that there is a disjunct q' in $q_{\mathcal{T}}$ that contains $R(x, y)$ and x is a root variable and $x \neq y$. We can assume the following:

- (A) there is no $q'' \subsetneq q'$ with $q'' \sqsubseteq_{\mathcal{T}} q$.
- (B) $(q' \setminus \{S(x, z)\}) \cup \{S(x, z')\} \not\sqsubseteq_{\mathcal{T}} q$ where
 - x is root variable in q' ,
 - z occurs more than once, and
 - z' is a fresh variable.

If Point (A) is not satisfied, set $q'_{\mathcal{T}} = q_{\mathcal{T}} \vee q'$. It is clear that $q'_{\mathcal{T}} \equiv q_{\mathcal{T}}$ and that during reduction of $q'_{\mathcal{T}}$ the disjunct q_i will be removed since $q_i \sqsubseteq q'_i$. However, q'_i will not be removed since q_i was also not removed from $q_{\mathcal{T}}$. If Point (B) is not satisfied, the same arguments apply.

We show that R is \mathcal{T} -generated in q . Since $q' \sqsubseteq_{\mathcal{T}} q$, Lemma 1 implies the existence of a homomorphism h from q to $\mathcal{I}_{q', \mathcal{T}}$. We distinguish two cases:

- Assume that q' is of the form $\{S_1(x, y), \dots, S_k(x, y)\}$, where the S_i are role names or inverse roles and R is among them. Hence, both x and y are root variables.
 - If q has two root variables and $k > 1$, it is clear that $q \equiv q'$, because of Points (A) and (B). Hence, R is \mathcal{T} -generated in q . If $k = 1$, then $\{R(x, y)\} = q' \sqsubseteq_{\mathcal{T}} q$.
 - If q has one root variable x_r , assume first that $h(x_r) \notin \{x, y\}$, i.e., x_r is mapped to some anonymous element of $\mathcal{I}_{q_1, \mathcal{T}}$. This clearly gives a contradiction to Point (B) since q is a simple tree query and, thus, not both x and y are in the image of h . Hence, we have $h(x_r) \in \{x, y\}$. Assume w.l.o.g. that $h(x_r) = x$. By Point (A), none of the atoms $S_i(x, y)$ can be dropped from q' . By Point (B), there has to be *one* variable $z \in \text{var}(q)$ such that $h(z) = y$ and $S_i(x_r, z) \in q$ for all $i \in \{1, \dots, k\}$. Hence we obtain $q \sqsubseteq q'$. By Lemma 2, q is a disjunct in $q_{\mathcal{T}}$. Since $q_{\mathcal{T}}$ is reduced, this implies $q \equiv q'$. Thus, q has two root variables, contradiction.

- Assume that q' contains an atom \mathbf{at}^* that involves x but not y . Thus, x is a root variable in q' and y is not. We distinguish cases regarding the root variable(s) of q . In the first two cases we identify atoms \mathbf{at} that can be dropped from q' such that still $q' \setminus \{\mathbf{at}\} \sqsubseteq_{\mathcal{T}} q$. However, this contradicts Point (A).
 - q has two root variables, i.e., it is of the form $\{S_1(x_r, y_r), \dots, S_k(x_r, y_r)\}$. Let first be $k > 1$. If $(h(x_r), h(y_r)) \in \{(x, y), (y, x)\}$, then \mathbf{at}^* can be dropped. Otherwise, we can drop $R(x, y)$.
If $k = 1$ and $(h(x_r), h(y_r)) \in \{(x, y), (y, x)\}$, then \mathbf{at}^* can be dropped. Otherwise, if either $h(x_r)$ or $h(y_r)$ are in the anonymous part of $\mathcal{I}_{q', \mathcal{T}}$, then one can drop all but one atoms from q' , i.e., either \mathbf{at}^* or $R(x, y)$. If $S_1(h(x_r), h(y_r)) \in q'$, one can drop all atoms except for $S_1(h(x_r), h(y_r))$.
 - q has precisely one root variable x_r and h maps x_r to $y' \neq x$. If $y' = y \cdot c_{R_1} \cdots c_{R_n}$ (possibly $n = 0$), then \mathbf{at}^* can be dropped from q_1 , since q is a simple tree query. Analogously, if $y' \neq y \cdot c_{R_1} \cdots c_{R_n}$, one can drop all but one atoms from q' , i.e., one of \mathbf{at}^* or $R(x, y)$ can be omitted.
 - The homomorphism h maps the root variable x_r of q to x . If there is $R(x_r, z) \in q$ such that $h(z) = y$, then $z \neq x_r$, and thus R is \mathcal{T} -generated in q . If not, there has to be either an atom $B(x_r) \in q$ and $\mathcal{T} \models \exists R \sqsubseteq B$ or an atom $S(x_r, z) \in q$ such that $z \neq x_r$ occurs only once and $\mathcal{T} \models \exists R \sqsubseteq \exists S$, otherwise Point (A) from above is not satisfied, as $R(x, y)$ can be dropped from q' . In both cases R is \mathcal{T} -generated in q . □

We are finally ready to prove Theorem 5 which we state again for completeness.

Theorem 5 Let \mathcal{T} be a DL-Lite-TBox. A \mathcal{T} -minimal CQ q is in PTIME relative to \mathcal{T} iff

1. q is a simple tree query, and
2. if r and r^- are \mathcal{T} -generated in q , then $\{r(x, y)\} \sqsubseteq_{\mathcal{T}} q$ or q is of the form $\{S_1(x, y), \dots, S_n(x, y)\}$.

Otherwise, q is #P-hard relative to \mathcal{T} .

Proof. “ \Rightarrow ”: We show the contrapositive. Assume first that q is not a simple tree query. Then Theorem 4 implies that answering q relative to \mathcal{T} is #P-hard. Assume now that q is a simple tree query but condition 2 from Theorem 5 is not satisfied, i.e., there is a role r such that both r and r^- can be generated in q w.r.t. \mathcal{T} but (i) $\{r(x, y)\} \not\sqsubseteq_{\mathcal{T}} q$ and (ii) $q \neq \{S_1(x, y), \dots, S_n(x, y)\}$. By Point 1 of Lemma 4, we can view $q_{\mathcal{T}}$ as a query where one variable x_r occurs in all atoms. We make a case distinction on how r and r^- are generated.

- $\{r(x_r, y), r(y, x_r)\} \subseteq q$. W.l.o.g. we can assume that q has to contain an atom \mathbf{at} that involves x_r but not y because of (ii). We assume here that $\mathbf{at} = A(x_r)$, the case where $\mathbf{at} = s(x_r, y')$ is treated analogously. By Lemma 2, we can assume that q is contained in $q_{\mathcal{T}}$. Since x_r is root variable of q but no separator, the same holds for $q_{\mathcal{T}}$. Because $q_{\mathcal{T}}$ is not ranked, Corollary 1 is not

applicable. However, due to the construction in (the proof of) Proposition 1, the ranking $\overline{q_{\mathcal{T}}}$ of $q_{\mathcal{T}}$ contains strict reducts q_1, q_2 of q such that

$$\begin{aligned} \{r^{<}(x_r, y_1), r^{>}(x_r, y_1), x_r < y_1, A(x_r)\} &\subseteq q_1 \\ \{r^{>}(y_2, x_r), r^{<}(y_2, x_r), x_r > y_2, A(x_r)\} &\subseteq q_2 \end{aligned}$$

According to Lemma 3, q_1 and q_2 are not removed during the reduction of $\overline{q_{\mathcal{T}}}$. Moreover, it is clear that $q_1 \not\equiv q_2$. Observe now that x_r is still root variable but no separator, e.g., in $r^{<}$ it occurs both in position 1 and 2. Now, the application of Corollary 1 to (the reduction of) $\overline{q_{\mathcal{T}}}$ yields #P-hardness of q .

- Assume $\{r(x_r, y), r(y, x_r)\} \not\subseteq q$. Since both r and r^- are \mathcal{T} -generated, Point 3 of Lemma 4 implies the existence of disjuncts q_1, q_2 in $q_{\mathcal{T}}$ such that

- q_1 contains $r(x_r, y_1)$ and $y_1 \neq x$, and
- q_2 contains $r(y_2, x_r)$ and $y_2 \neq x$.

By definition of $q_{\mathcal{T}}$ we have that $q_i \sqsubseteq_{\mathcal{T}} q$, hence $q_1 \neq \{r(x_r, y_1)\}$ and $q_2 \neq \{r(y_2, x_r)\}$ by (i). We argue first that y_1 is not a root variable of q_1 : If it is, q_1 is of the form $\{R_1(x_r, y_1), \dots, R_m(x_r, y_1)\}$ with $m > 1$ because of the previous argument. Since $q_1 \sqsubseteq_{\mathcal{T}} q$, there is a homomorphism h from q to $\mathcal{I}_{q_1, \mathcal{T}}$. W.l.o.g. we can assume that q_1 satisfies Points (A) and (B) from the proof of Lemma 4. If $m > 1$, then h has to ‘hit’ all atoms $R_1(x_r, y_1), \dots, R_m(x_r, y_1)$. This yields $q_1 \sqsubseteq q$ and thus $q_1 \equiv q$, since $q_{\mathcal{T}}$ is reduced. But this is in contradiction with (ii). The same arguments apply to variable y_2 in q_2 .

It is clear that x_r is not a separator because it occurs in position 1 and 2 for r . If $q_{\mathcal{T}}$ is ranked then it is #P-hard, by Corollary 1. If $q_{\mathcal{T}}$ is not ranked, consider its ranking $\overline{q_{\mathcal{T}}}$. The construction in (the proof of) Proposition 1 together with Lemma 3 imply the existence of disjuncts q'_1, q'_2 such that

- $\{r^{<}(x_r, y_1), x < y_1\} \subseteq q'_1$ and
- $\{r^{<}(y_2, x_r), y_2 < x\} \subseteq q'_2$

where in both queries x_r is a root variable but y_1 and y_2 are not. Thus, $\overline{q_{\mathcal{T}}}$ does not have a separator and Corollary 1 implies #P-hardness of $\overline{q_{\mathcal{T}}}$ and thus of q .

“ \Leftarrow ”. Assume that q and \mathcal{T} satisfy Conditions 1 and 2 from the Theorem. By Theorem 1, it suffices to show that, for the FO-rewriting $q_{\mathcal{T}}$ of q relative to \mathcal{T} , answer probabilities over tuple-independent databases can be computed in PTIME. Hence, let D be the ipABox \mathcal{A} viewed as a tuple-independent database. For a Boolean CQ p , we write $p_D(p)$ instead of $p_D^d((\cdot) \in p)$ and omit D whenever it is clear from the context.

By Point 2 of Lemma 4, we may assume that $q_{\mathcal{T}}$ is a UCQ in which all disjuncts are simple tree queries. Hence, we can assume w.l.o.g. that there is a variable x_r occurring in all atoms of $q_{\mathcal{T}}$. Lemma 4 helps us also to prove the following useful claim.

Claim. Let r be a role name such that either r or r^- is not \mathcal{T} -generated in q . There are no two disjuncts q_1 and q_2 of $q_{\mathcal{T}}$ such that $r(x_r, y_1) \in q_1$ with $x_r \neq y_1$ and $r(y_2, x_r) \in q_2$ with $x_r \neq y_2$.

Proof of Claim. Assume to the contrary that there are such q_1, q_2 . By Point 3 of Lemma 4, we have that both r and r^- are \mathcal{T} -generated in q or $\{r(x, y)\} \sqsubseteq_{\mathcal{T}} q$. The former is excluded by our assumption, hence the latter is the case. Set $q'_{\mathcal{T}} = q_{\mathcal{T}} \vee \{r(x, y)\}$. It is easy to see that $q'_{\mathcal{T}} \equiv q_{\mathcal{T}}$ and that during reduction of $q'_{\mathcal{T}}$ the added disjunct is not removed but q_1 and q_2 are (unless they are equivalent to $\{r(x, y)\}$). This finishes the proof of the claim.

Now, we distinguish three cases.

Case (i). For all role names r , either r or r^- is not \mathcal{T} -generated in q . The above claim provides a necessary condition for x_r being a separator. However, it is not yet sufficient, since there might be ‘reflexive’ atoms of the form $r(x_r, x_r)$ occurring in $q_{\mathcal{T}}$. We show that these atoms can be removed in the following way. If \sim is an equivalence relation over $\text{var}(q')$ and at an atom, we define at_{\sim} by taking

$$A(x)_{\sim} = A(x)$$

$$R(x, y)_{\sim} = \begin{cases} R^{\neq}(x, y) & \text{if } x \approx y \\ R^=(x) & \text{if } x \sim y \end{cases}$$

Denote with $\text{EQ}(q')$ the set of all equivalence relations over $\text{var}(q')$. Then, define $q'_{\mathcal{T}}$ as

$$\bigvee_{q' \in q_{\mathcal{T}}} \bigvee_{\text{EQ}(q')} \{\text{at}_{\sim} \mid \text{at} \in q'\}$$

It is not hard to prove that D can be transformed in polynomial time into some D' such that $P_D(q_{\mathcal{T}}) = P_{D'}(q'_{\mathcal{T}})$. Moreover, it should be clear that the statement of the claim remains true after this transformation. Finally observe that x_r is a separator variable in $q'_{\mathcal{T}}$.

We now perform the following steps in order to compute $p_{D'}(q'_{\mathcal{T}})$:

- If $q'_{\mathcal{T}}$ is not symbol-connected and its symbol-connected components are q_1, \dots, q_k , compute:

$$P(q'_{\mathcal{T}}) = 1 - (1 - P(q_1)) \cdot \dots \cdot (1 - P(q_k))$$

- Continue with q_i , which is clearly symbol-connected and, by construction, still has the separator x_r . Compute

$$P(q_i) = 1 - \prod_{a \in \text{Ind}(\mathcal{A})} (1 - P(q_i[a/x_r]))$$

- The query $q_i[a/x_r]$ comprises disjuncts q'_i with $\text{var}(q'_i) = \{z_1, \dots, z_m\}$ of the form

$$A_1(a) \wedge \dots \wedge A_n(a) \wedge q'_{i1} \wedge \dots \wedge q'_{im} \quad (\star)$$

where each q'_{ij} is obtained from q'_i by restricting to all atoms that involve z_j , i.e., q'_{ij} is of the form $S_1(a, z_j), \dots, S_\ell(a, z_j)$.

In order to compute $p(q_i[a/x_r])$, we apply the distributivity law and distribute the ‘ \wedge ’s in (\star) over the (outer) disjunctions. We obtain a conjunction

of N disjunctive sentences $q_\wedge = p_1 \wedge \dots \wedge p_N$. The restricted form of the conjuncts in (\star) implies that each p_i is a disjunction of CQs of the form $A(a)$ or $S_1(a, y), \dots, S_\ell(a, y)$.

- The probability $p(q_\wedge)$ can be computed using the (dual) inclusion/exclusion principle as follows:

$$p(q_\wedge) = \sum_{\emptyset \subset Y \subseteq [N]} (-1)^{|Y|+1} P(q_Y)$$

where $[N] = \{1, \dots, N\}$ and $q_Y = \bigvee_{y \in Y} p_y$. Note that q_Y is a disjunctive sentence for all $Y \subseteq [N]$.

- Let us analyze the structure of q_Y for some (non-empty) $Y \subseteq [N]$. Since its disjuncts are obtained from some combination of the p_i , q_Y can be partitioned into q_{A_1}, \dots, q_{A_k} , and $q_{Y'}$ such that
 - $q_{A_i} = A_i(a)$ and
 - $q_{Y'}$ contains only CQs of the form $S_1(a, y), \dots, S_\ell(a, y)$

In tuple-independent databases, the disjuncts of the former form are pairwise independent of each other and independent of disjuncts of the latter form, so compute

$$p(q_Y) = 1 - (1 - p(q_{Y'})) \cdot \prod_{i=1}^k (1 - p(q_{A_i}))$$

- Observe that $p(q_{A_i})$ can be read off from D for all i .
- For computing $p(q_{Y'})$ observe that $q_{Y'}$ can be viewed as a query with a separator z , since every atom has precisely one free variable. Hence, we can compute

$$p(q_{Y'}) = 1 - \prod_{b \in \text{Ind}(\mathcal{A})} (1 - p(q_{Y'}[b/z]))$$

- The query $q_{Y'}[b/z]$ consists of disjuncts of the form $S_1(a, b) \wedge \dots \wedge S_\ell(a, b)$. Again, distribute the \wedge 's over the outer disjunctions to obtain a conjunction of disjunctive sentences where each disjunctive sentence comprises only *one* atom $S(a, b)$ and apply the (dual) inclusion/exclusion principle. Each disjunctive sentence q_d obtained in this way is of the form $S_1(a, b) \vee \dots \vee S_n(a, b)$. Hence, its probability can be computed as

$$p(q_d) = 1 - (1 - p(S_1(a, b))) \cdots (1 - p(S_n(a, b)))$$

which can be read off from D .

Case (ii). There is a role name r such that both r and r^- is \mathcal{T} -generated in q , but q is not of the form $\{S_1(x, y), \dots, S_k(x, y)\}$. Statement 2 from the Theorem tells us that for those r we have $\{r(x, y)\} \sqsubseteq_{\mathcal{T}} q$. In the same way as in the proof of the claim we obtain that there are no disjuncts q_1 and q_2 in $q_{\mathcal{T}}$ with the properties described in the statement of the claim. Hence, we can continue with computing $p_D(q_{\mathcal{T}})$ in the same way as in Case (i).

Case (iii). There is a role name r such that both r and r^- is \mathcal{T} -generated in q and q is of the form $\{S_1(x, y), \dots, S_k(x, y)\}$. It is not hard to prove that $q' \sqsubseteq_{\mathcal{T}} q$ implies actually $q' \sqsubseteq q$. Hence, $q_{\mathcal{T}} \equiv q$ and it suffices to compute $p_D(q)$. Again, x_r is root variable but no separator, since q contains atoms $r(x, y)$ and $r(y, x)$ (recall that r and r^- are \mathcal{T} -generated). Thus, q is not ranked and we need to consider the ranking \bar{q} of q in order to compute $p(q)$. Following (the proof of) Proposition 1, \bar{q} is obtained as

$$\begin{aligned} & S_1^=(x), \dots, S_k^=(x) \\ & \vee S_1^<(x, y), \dots, S_k^<(x, y), x < y \\ & \vee S_1^>(x, y), \dots, S_k^>(x, y), x > y \end{aligned}$$

where $S_i^= = s_i^-$ (independent from S_i being s_i or s_i^-), $S_i^<$ is $s_i^<$ ($s_i^>$, respectively) when $S_i = s_i$ ($S_i = s_i^-$, respectively), and analogously for $S_i^>$. It is easy to see that \bar{q} has the separator x . Observe that the ranking also removes reflexive atoms, hence we can continue with computing $p(\bar{q})$ using the algorithm in Case (i). \square

B Proofs for Section 5

Before we can prove Theorem 6, we need to prove some intermediate lemmas. In particular, the central observation is as follows.

Lemma 5. *If a Boolean rooted CQ q is not FO-rewritable relative to an \mathcal{ELI} -TBox \mathcal{T} , then there exists an ABox \mathcal{A} and assertions $R_3(a_3, a_2)$, $R_2(a_2, a_1)$, $R_1(a_1, a_0)$ such that $\mathcal{A}, \mathcal{T} \models q$, but $\mathcal{A}', \mathcal{T} \not\models q$ when \mathcal{A}' is \mathcal{A} with any of the assertions $R_3(a_3, a_2)$, $R_2(a_2, a_1)$, $R_1(a_1, a_0)$ dropped.*

The proof consists of a sequence of ABox transformations and uses a notion of boundedness similar to the one introduced in [27], but adapted from instance queries to CQs.

Let \mathcal{T} be an \mathcal{ELI} -TBox and q a Boolean connected CQ that involves at least one individual name. We first show that we can assume w.l.o.g. that \mathcal{T} contains only CIs of the forms

$$\begin{aligned} A &\sqsubseteq B & A &\sqsubseteq \exists R.B \\ B_1 \sqcap B_2 &\sqsubseteq A & \exists R.B &\sqsubseteq A \end{aligned}$$

where R ranges over role names and their inverse and A, B, B_1, B_2 over concept names and \top . Let $\text{sub}(\mathcal{T})$ denote the set of all subconcepts of (concepts that occur in) \mathcal{T} and reserve a concept name X_C for every $C \in \text{sub}(\mathcal{T}) \setminus (\mathbf{N}_C \cup \{\top\})$ such that X_C occurs neither in \mathcal{T} nor in q . Set

$$\sigma(C) = \begin{cases} C & \text{if } C \in \mathbf{N}_C \cup \{\top\} \\ X_{D_1} \sqcap X_{D_2} & \text{if } C = D_1 \sqcap D_2 \\ \exists r.X_D & \text{if } C = \exists r.D \end{cases}$$

Then put

$$\mathcal{T}' = \bigcup_{C \sqsubseteq D \in \mathcal{T}} X_C \sqsubseteq X_D \cup \bigcup_{C \in \text{sub}(\mathcal{T}) \setminus (\mathcal{N}_c \cup \{\top\})} X_C \equiv \sigma(C)$$

where $C \equiv D$ abbreviates $C \sqsubseteq D, D \sqsubseteq C$. After further replacing each CI of the form $A \sqsubseteq B_1 \sqcap B_2$ with $A \sqsubseteq B_1$ and $A \sqsubseteq B_2$, \mathcal{T}' is of the required syntactic form. Clearly, the conversion can be done in polynomial time.

We want to replace \mathcal{T} with the TBox \mathcal{T}' in normal form. To implement this, we consider ABoxes in a restricted signature. A *predicate* is a set of concept names and role names and a *signature* is a set of predicates. We use $\text{sig}(\mathcal{T})$ to denote the set of predicates that occur in \mathcal{T} and likewise for $\text{sig}(q)$. A Σ -ABox is an ABox that contains only symbols from Σ . Due to the following result, we shall indeed be able to replace \mathcal{T} with \mathcal{T}' when we are careful about ABox signatures.

Theorem 10. *Let $\Sigma = \text{sig}(\mathcal{T}) \cup \text{sig}(q)$.*

1. *q is FO-rewritable relative to \mathcal{T} (over all ABoxes) iff q is FO-rewritable relative to \mathcal{T}' over Σ -ABoxes;*
2. *q is #P-hard relative to \mathcal{T} (over all ipABoxes) iff q is #P-hard relative to \mathcal{T}' over Σ -ipABoxes.*

Proof. For Point 1, first assume that q is FO-rewritable relative to \mathcal{T} and let $q_{\mathcal{T}}$ be an FO-rewriting. We show that $q_{\mathcal{T}}$ is also an FO-rewriting of q relative to \mathcal{T}' over Σ -ABoxes. To this end, let \mathcal{A} be a Σ -ABox. Since the fresh concept names X_C occur neither in q nor in \mathcal{A} , it is easy to show that $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{A}, \mathcal{T}' \models q$. In summary, $\mathcal{A}, \mathcal{T}' \models q$ iff $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{I}_{\mathcal{A}} \models q_{\mathcal{T}}$, thus we are done. Conversely, assume that q is FO-rewritable relative to \mathcal{T}' over Σ -ABoxes and let $q_{\mathcal{T}'}$ be an FO-rewriting. Since each non- Σ -symbol is interpreted as the empty set in $\mathcal{I}_{\mathcal{A}}$ for any Σ -ABox \mathcal{A} , we can w.l.o.g. assume that no such symbol occurs in $q_{\mathcal{T}'}$ (if it does, replace it with **false**). We show that $q_{\mathcal{T}'}$ is also an FO-rewriting of q relative to \mathcal{T} . Let \mathcal{A} be an ABox and $\mathcal{A}|_{\Sigma}$ the result of dropping all non- Σ -assertions from \mathcal{A} . We have $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{A}|_{\Sigma}, \mathcal{T} \models q$ (since q and \mathcal{T} contain only Σ -symbols) iff $\mathcal{A}|_{\Sigma}, \mathcal{T}' \models q$ (since the X_C occur neither in q nor in $\mathcal{A}|_{\Sigma}$) iff $\mathcal{I}_{\mathcal{A}|_{\Sigma}} \models q_{\mathcal{T}'}$ (since $q_{\mathcal{T}'}$ is an FO-rewriting) iff $\mathcal{I}_{\mathcal{A}} \models q_{\mathcal{T}'}$ (since there are no non- Σ -symbols in $q_{\mathcal{T}'}$).

For Point 2, first assume that q is #P-hard relative to \mathcal{T} . Since \mathcal{T} and q contain only Σ -symbols, we have $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{A}|_{\Sigma}, \mathcal{T} \models q$ for all ABoxes \mathcal{A} . It follows that q is #P-hard relative to \mathcal{T} over Σ -ABoxes. Since we have $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{A}, \mathcal{T}' \models q$ for all Σ -ABoxes \mathcal{A} (see proof of Point 1 above) and thus also $p(\mathcal{A}, \mathcal{T} \models q) = p(\mathcal{A}, \mathcal{T}' \models q)$, answering q relative to \mathcal{T} over Σ -ABoxes is simply the same problem as answering q relative to \mathcal{T}' over Σ -ABoxes and we are done. For the converse direction, $p(\mathcal{A}, \mathcal{T} \models q) = p(\mathcal{A}, \mathcal{T}' \models q)$ for all Σ -ABoxes \mathcal{A} means that answering q relative to \mathcal{T}' over Σ -ABoxes is simply a subproblem of answering q relative to \mathcal{T} , thus #P-hardness of the former implies #P-hardness of the latter. \square

The proof of Theorem 6 is based on a connection between FO-rewritability and boundedness of an appropriate fixpoint operator, similar to what was observed in [27]. Given an ABox \mathcal{A} and an $a \in \text{Ind}(\mathcal{A})$, we denote with $\mathcal{A}|_a$ the *neighbourhood* of a , i.e., the restriction of \mathcal{A} to the individual name a and all members of $\{b \mid R(a, b) \in \mathcal{A}\} \cup \{b \mid R(b, a) \in \mathcal{A}\}$ (where R is a role name or its inverse). For a TBox \mathcal{T} , define

$$f_{\mathcal{T}}(\mathcal{A}) = \mathcal{A} \cup \{A(a) \mid a \in \text{Ind}(\mathcal{A}) \wedge \mathcal{A}|_a, \mathcal{T} \models A(a)\}$$

and set $f_{\mathcal{T}}^{\infty}(\mathcal{A}) := \bigcup_{i \geq 0} f_{\mathcal{T}}^i(\mathcal{A})$, where $f_{\mathcal{T}}^i(\cdot)$ denotes application of $f_{\mathcal{T}}$, iterated i times. Note that the application of $f_{\mathcal{T}}(\cdot)$ yields ABoxes, though not necessarily Σ -ABoxes, $\Sigma = \text{sig}(\mathcal{T}) \cup \text{sig}(q)$. It is not hard to prove that for all $A \in \mathbf{N}_{\mathcal{C}}$ and $a \in \mathbf{N}_{\mathcal{I}}$, we have $\mathcal{A}, \mathcal{T} \models A(a)$ iff $A(a) \in f_{\mathcal{T}}^{\infty}(\mathcal{A})$ [27]. We want to establish an analogous claim for CQs, which requires first introducing several technical notions. In particular, we construct a query $q_{\mathcal{T}}$ from q and \mathcal{T} such that answering q over a Σ -ABox \mathcal{A} is equivalent to answering $q_{\mathcal{T}}$ over $f_{\mathcal{T}}^{\infty}(\mathcal{A})$. Note that $q_{\mathcal{T}}$ is in general *not* an FO-rewriting of q relative to \mathcal{T} , which are not guaranteed to exist in \mathcal{ELI} ; intuitively, $q_{\mathcal{T}}$ constitutes the part of the FO-rewriting that deals with objects generated by existential quantifiers.

We use $\text{EQ}(q)$ to denote the set of all equivalence relations on $\text{term}(q)$ such that $a \not\sim b$ for all distinct $a, b \in \mathbf{N}_{\mathcal{I}}$. Every $\sim \in \text{EQ}(q)$ defines the following *collapsing* of q :

$$q_{\sim} = \{r(s_{[t]}, s_{[t']}) \mid r(t, t') \in q\} \cup \{A(s_{[t]}) \mid A(t) \in q\}.$$

where $s_{[t]} = a$ if the individual name a is in $[t]$ and $s_{[t]}$ is the fresh variable $z_{[t]}$ otherwise.

We call a CQ *tree-shaped* if the undirected graph (V_q, E_q) with $V_q = \text{term}(q)$ and $E_q = \{\{t, t'\} \mid r(t, t') \in q\}$ is a tree in which an individual name occurs, if at all, only at the root. A *splitting* S of a CQ q is a partition q_0, \dots, q_k of q (with q_0 possibly empty and q_1, \dots, q_k non-empty) such that for $1 \leq i \leq k$, we have

1. q_1, \dots, q_k are tree-shaped queries with roots t_1, \dots, t_k
2. $\text{term}(q_i) \cap \text{term}(q_j) = \emptyset$ for $i < j \leq k$
3. $\text{term}(q_0) \cap \text{term}(q_i) = \{t_i\}$.

Let $\text{split}(q)$ denote the set of all splittings of q and $\text{CN}(\mathcal{T})$ the concept names that occur in \mathcal{T} . For a subset $\rho \subseteq \text{CN}(\mathcal{T})$, we use \mathcal{A}_{ρ} to denote the ABox $\{A(\hat{a}) \mid A \in \rho\}$ where \hat{a} is a fixed, ‘special’ individual name. Given a splitting $S = q_0, \dots, q_k$ of a CQ q , a map $\rho : \{1, \dots, k\} \rightarrow 2^{\text{CN}(\mathcal{T})}$ is a *justification for S relative to \mathcal{T}* if for all $i \in \{1, \dots, k\}$, we have $\mathcal{A}_{\rho(i)}, \mathcal{T} \models^{\hat{a}} q_i$ meaning that in every model \mathcal{I} of $\mathcal{A}_{\rho(i)}$ and \mathcal{T} , there is a match of the tree-shaped Boolean query q_i that maps its root t_i to $\hat{a}^{\mathcal{I}}$. We use $q_{\rho(i)}$ to denote the query $\bigwedge_{A \in \rho(i)} A(t_i)$ where, again t_i denotes the root of q_i . Let $\text{just}(S, \mathcal{T})$ denote all possible justifications of S relative to \mathcal{T} . For each $\sim \in \text{EQ}(q)$, set

$$\hat{q}_{\sim} = \exists z_{[\hat{t}_1]} \cdots \exists z_{[\hat{t}_n]} \bigvee_{S=q_0, \dots, q_k \in \text{split}(q_{\sim})} \bigvee_{\rho \in \text{just}(S, \mathcal{T})} (q_0 \wedge \bigwedge_{i=1}^k q_{\rho(i)})$$

where $[\hat{t}_1], \dots, [\hat{t}_n]$ are the equivalence classes of \sim that do not contain an individual name. Finally, we define the CQ $q_{\mathcal{T}}$ as $\bigvee_{\sim \in \text{EQ}(q)} \hat{q}_{\sim}$. Clearly, $q_{\mathcal{T}}$ is a UCQ.

Theorem 11. *For any CQ q , \mathcal{ELL} -TBox \mathcal{T} in normal form, and Σ -ABox \mathcal{A} , where $\Sigma = \text{sig}(\mathcal{T}) \cup \text{sig}(q)$, we have $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{I}_{f_{\hat{\sim}}(\mathcal{A})} \models q_{\mathcal{T}}$.*

To prove Theorem 11, we introduce canonical models for ABoxes and \mathcal{ELL} -TBoxes. Let \mathcal{A} be an ABox and \mathcal{T} an \mathcal{ELL} -TBox in normal form. A *type for \mathcal{T}* is a set $T \subseteq \text{CN}(\mathcal{T})$. When $a \in \text{Ind}(\mathcal{A})$, T, T' are types for \mathcal{T} , and R is a role, we write

- $a \rightsquigarrow_R T$ if $\mathcal{A}, \mathcal{T} \models \exists R. \Box T(a)$ and for each type S for \mathcal{T} with $T \subsetneq S$, we have $\mathcal{A}, \mathcal{T} \not\models \exists R. \Box T(a)$;
- $T \rightsquigarrow_R T'$ if $\mathcal{T} \models \Box T \sqsubseteq \exists R. \Box T'$ and for each type S for \mathcal{T} with $T' \subsetneq S$, we have $\mathcal{T} \not\models \Box T \sqsubseteq \exists R. \Box S$.

A *path for \mathcal{A} and \mathcal{T}* is a sequence $p = aR_1T_1 \cdots T_{n-1}R_nT_n$, $n \geq 0$, with $a \in \text{Ind}(\mathcal{A})$, R_1, \dots, R_n roles, and T_1, \dots, T_n types for \mathcal{T} such that $a \rightsquigarrow_{R_1} T_1$ and $T_i \rightsquigarrow_{R_i} T_{i+1}$ for $1 \leq i < n$. When $n > 0$, we use $\text{tail}(p)$ to denote T_n . Define the interpretation $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ as follows:

$$\begin{aligned} \Delta^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \text{the set of all paths for } \mathcal{A} \text{ and } \mathcal{T} \\ A^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \{a \in \text{Ind}(\mathcal{A}) \mid \mathcal{A}, \mathcal{T} \models A(a)\} \cup \\ &\quad \{p \in \Delta^{\mathcal{I}} \setminus \text{Ind}(\mathcal{A}) \mid \mathcal{T} \models \Box \text{tail}(p) \sqsubseteq A\} \\ r^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= \{(a, b) \in \text{Ind}(\mathcal{A}) \times \text{Ind}(\mathcal{A}) \mid r(a, b) \in \mathcal{A}\} \cup \\ &\quad \{(p, prT) \mid prT \in \text{Paths}\} \cup \\ &\quad \{(pr^-T, p) \mid pr^-T \in \text{Paths}\} \\ a^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}} &= a \end{aligned}$$

It is standard to prove that $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ is canonical in the following sense.

Lemma 6. *For any CQ q , \mathcal{ELL} -TBox \mathcal{T} in normal form, and ABox \mathcal{A} , we have $\mathcal{A}, \mathcal{T} \models q$ iff $\mathcal{I}_{\mathcal{A}, \mathcal{T}} \models q$.*

We are now ready to prove Theorem 11.

Proof.(sketch of Theorem 11) “if”. Assume that $\mathcal{I}_{f_{\hat{\sim}}(\mathcal{A})} \models q_{\mathcal{T}}$ and let π be a match of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\hat{\sim}}(\mathcal{A})}$. Let q' be the disjunct of $q_{\mathcal{T}}$ with $\mathcal{I}_{f_{\hat{\sim}}(\mathcal{A})} \models q'$. By definition of $q_{\mathcal{T}}$, there is an $S = q_0, \dots, q_k \in \text{split}(q_{\sim})$ and a $\rho \in \text{just}(S, \mathcal{T})$ such that $q' = q_0 \wedge \bigwedge_{i=1}^k q_{\rho(i)}$. For $1 \leq i \leq k$, we have $\mathcal{A}_{\rho(i)}, \mathcal{T} \models^{\hat{a}} q_i$ and, by Lemma 6, there is thus a match π_i of q_i in $\mathcal{I}_{\mathcal{A}_{\rho(i)}, \mathcal{T}}$ such that the root t_i of q_i is mapped to \hat{a} . Since π is a match of $q_{\rho(i)}$ in $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$, it can be shown that there is a homomorphism from $\mathcal{I}_{\mathcal{A}_{\rho(i)}, \mathcal{T}}$ to $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ with $h(\hat{a}) = \pi(t_i)$. Thus, we find a match π'_i of q_i in $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ with $\pi'_i(t_i) = \pi(t_i)$. It can be verified that $\pi \cup \pi_1 \cup \dots \cup \pi_k$ is a match of q in $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$. By Lemma 6, we have $\mathcal{A}, \mathcal{T} \models q$ as required.

“only if”. Assume that $\mathcal{A}, \mathcal{T} \models q$. By Lemma 6, this means $\mathcal{I}_{\mathcal{A}, \mathcal{T}} \models q$. Let π be a match of q in $\mathcal{I}_{\mathcal{A}, \mathcal{T}}$ and let a_1, \dots, a_k be the elements of $\text{Ind}(\mathcal{A})$ that are in the range of π . Define a splitting $S = q_0, \dots, q_k$ of q , where q_0 is the restriction

of q to the terms t with $\pi(t) \in \text{Ind}(\mathcal{A})$ and each q_i is the restriction of q to the terms t with $\pi(t) = a_i p$ for some non-empty p . Define a justification ρ for S relative to \mathcal{T} by setting $\rho(i) = \{A \in \text{CN}(\mathcal{T}) \mid a_i \in A^{\mathcal{I}_{\mathcal{A}, \mathcal{T}}}\}$ for $1 \leq i \leq k$. Then $q' = q_0 \wedge \bigwedge_{i=1}^k q_{\rho(i)}$ is a disjunct in $q_{\mathcal{T}}$ and it can be verified that $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})} \models q'$. \square

We say that a CQ q is *k-bounded relative to a TBox \mathcal{T} over Σ -ABoxes* if for every Σ -ABox \mathcal{A} , we have that $\mathcal{I}_{f_{\mathcal{T}}^k(\mathcal{A})} \models q_{\mathcal{T}}$ iff $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})} \models q_{\mathcal{T}}$. We say that q is *bounded relative to a TBox \mathcal{T} over Σ -ABoxes* if it is k -bounded for some k .

Theorem 12. *If a CQ q is bounded relative to \mathcal{T} over Σ -ABoxes, $\Sigma = \text{sig}(\mathcal{T}) \cup \text{sig}(q)$, then it is FO-rewritable relative to \mathcal{T} over Σ -ABoxes.*

Proof. Assume that q is bounded relative to \mathcal{T} over Σ -ABoxes and let $k > 0$ be such that $\mathcal{I}_{f_{\mathcal{T}}^k(\mathcal{A})} \models q_{\mathcal{T}}$ iff $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})} \models q_{\mathcal{T}}$ for every Σ -ABox \mathcal{A} .

Observe that we work with finite Σ , and thus there are only finitely many neighborhoods \mathcal{A}_a that can occur in a Σ -ABox \mathcal{A} , up to isomorphism. Every such neighborhood \mathcal{N} with individual name a in the center can be converted in a straightforward way into an existential, conjunctive, positive FO-formula:

$$\varphi_{\mathcal{N}} = \bigwedge_{A(a) \in \mathcal{N}} A(x) \wedge \bigwedge_{b \in \text{Ind}(\mathcal{N})} \exists y : \left(\bigwedge_{R(a,b) \in \mathcal{N}} R(x,y) \wedge \bigwedge_{B(b) \in \mathcal{N}} B(y) \right)$$

where $R(x, y)$ denotes $r(y, x)$ when $R = r^-$. For each concept name A , we use Γ_A to denote the set of neighborhoods \mathcal{N} with center a such that $\mathcal{N} \models A(a)$. For every concept name A and $i \geq 0$, set

$$\begin{aligned} - q_A^0(x) &:= A(x) \\ - q_A^{i+1}(x) &:= q_A^i \vee \bigvee_{\mathcal{N} \in \Gamma_A} \varphi'_{\mathcal{N}} \text{ where } \varphi'_{\mathcal{N}} \text{ is obtained from } \varphi_{\mathcal{N}} \text{ by replacing, for} \\ &\quad \text{each concept name } B, \text{ every atom } B(z) \text{ with } q_B^i[z/x]. \end{aligned}$$

The following can be proved by induction on i :

Claim. For every Σ -ABox \mathcal{A} and $i \geq 0$, we have $\mathcal{I}_{\mathcal{A}} \models q_A^i[a]$ iff $A(a) \in f_{\mathcal{T}}^i(\mathcal{A})$.

Let $\widehat{q}_{\mathcal{T}}$ be $q_{\mathcal{T}}$ with every atom $A(t)$ replaced with $q_A^k[t/x]$. We show that $\widehat{q}_{\mathcal{T}}$ is an FO-rewriting of q relative to \mathcal{T} , which finishes the proof. By the above claim and due to the fact that $f_{\mathcal{T}}^k(\mathcal{A})$ differs from \mathcal{A} only by additional concept assertions, we have $\mathcal{I}_{\mathcal{A}} \models \widehat{q}_{\mathcal{T}}$ iff $\mathcal{I}_{f_{\mathcal{T}}^k(\mathcal{A})} \models q_{\mathcal{T}}$. This is the case iff $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})} \models q_{\mathcal{T}}$ (by choice of k) iff $\mathcal{A}, \mathcal{T} \models q$ (by Theorem 11). \square

The next and central step is to prove the following result.

Theorem 13. *If a CQ q is unbounded relative to an \mathcal{ELLI} -TBox \mathcal{T} over Σ -ABoxes, then q is $\#P$ -hard relative to \mathcal{T} over Σ -ABoxes.*

Let q be unbounded relative to \mathcal{T} over Σ -ABoxes and set

$$m = |\mathcal{T}| \cdot |q_{\mathcal{T}}| \cdot (|\mathcal{T}| + |q_{\mathcal{T}}|)^{|q_{\mathcal{T}}|+2} + 1.$$

Since q is not bounded relative to \mathcal{T} , there is a Σ -ABox \mathcal{A}_0 such that $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_0)} \models q_{\mathcal{T}}$, but $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_0)} \not\models q_{\mathcal{T}}$. Choose a concrete match π of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_0)}$. In the following, when writing $R(a, b) \in \mathcal{A}$ with $R = r^-$, we mean $r(b, a) \in \mathcal{A}$. Construct a Σ -ABox \mathcal{A}_1 as follows:

- a π -path in \mathcal{A}_0 of length n is a sequence $p = a_0 R_1 a_1 \cdots R_n a_n$ such that $a_0 \in \text{ran}(\pi)$ and $R_i(a_{i-1}, a_i) \in \mathcal{A}_0$ for $1 \leq i \leq n$; we use $\text{tail}(p)$ to denote a_n ;
- $\text{Ind}(\mathcal{A}_1)$ is the set of all π -paths in \mathcal{A}_0 of length at most $m + |q_{\mathcal{T}}|$;
- if $A(a) \in \mathcal{A}_0$, $p \in \text{Ind}(\mathcal{A}_1)$, and $\text{tail}(p) = a$, then $A(p) \in \mathcal{A}$;
- if $pRa \in \text{Ind}(\mathcal{A}_1)$, then $R(p, pRa) \in \mathcal{A}_1$;
- if $r(a, b) \in \mathcal{A}_0$ and $a, b \in \text{ran}(\pi)$, then $r(a, b) \in \mathcal{A}_1$.
- these are all assertions in \mathcal{A}_1 .

Note that \mathcal{A}_1 has, in a loose sense, a forest shape: the roots, which are the elements of $\text{Ind}(\mathcal{A}_1) \cap \text{Ind}(\mathcal{A}_0)$, form a ‘core’ whose relational structure is not restricted in any way; moreover, each root a gives rise to a tree-shaped sub-ABox of \mathcal{A}_1 , namely the restriction to the individuals p that have a as a prefix.

Lemma 7. $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_1)} \models q_{\mathcal{T}}$, but $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_1)} \not\models q_{\mathcal{T}}$.

Proof. We first prove the following.

Claim. For all $i \leq m$, $p \in \text{Ind}(\mathcal{A}_1)$ of length at most $m + |q_{\mathcal{T}}| - i$ with $\text{tail}(p) = a$, and $A \in \mathbf{N}_{\mathcal{C}}$, we have $A(p) \in f_{\mathcal{T}}^i(\mathcal{A}_1)$ iff $A(a) \in f_{\mathcal{T}}^i(\mathcal{A}_0)$.

The proof is by induction on i . The induction start is trivial by construction of \mathcal{A}_1 . For the induction step, we have $A(p) \in f_{\mathcal{T}}^i(\mathcal{A}_1)$ iff $f_{\mathcal{T}}^{i-1}(\mathcal{A}_1)|_a, \mathcal{T} \models A(p)$ (by definition of the function $f_{\mathcal{T}}$) iff $f_{\mathcal{T}}^{i-1}(\mathcal{A}_0)|_a, \mathcal{T} \models A(a)$ (by induction hypothesis) iff $A(a) \in f_{\mathcal{T}}^i(\mathcal{A}_0)$.

To see that $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_1)} \models q_{\mathcal{T}}$, consider the match π used in the construction of \mathcal{A}_1 . By the claim, for all $a \in \text{ran}(\pi)$ and $A \in \mathbf{N}_{\mathcal{C}}$, we have $A(a) \in f_{\mathcal{T}}^m(\mathcal{A}_0)$ iff $A(a) \in f_{\mathcal{T}}^m(\mathcal{A}_1)$. It follows that π is also a match of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_1)}$.

For $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_1)} \not\models q_{\mathcal{T}}$, assume that there is a match τ of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_1)}$ and select a concrete disjunct q' of $q_{\mathcal{T}}$ that is matched by τ . Since q is connected and contains at least one individual name, by construction the same holds for q' . It follows that all paths $p \in \text{ran}(\tau)$ are of length at most $|q_{\mathcal{T}}|$. By the claim, it follows that for all $p \in \text{ran}(\tau)$ with $\text{tail}(p) = a$ and $A \in \mathbf{N}_{\mathcal{C}}$, we have $A(a) \in f_{\mathcal{T}}^{m-1}(\mathcal{A}_0)$ iff $A(p) \in f_{\mathcal{T}}^{m-1}(\mathcal{A}_1)$. This, in turn, implies that the function τ' defined by setting $\tau'(t) := \text{tail}(\tau(t))$ for all $t \in \text{term}(q')$ is a match of q' in $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_0)}$, in contradiction to $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_0)} \not\models q_{\mathcal{T}}$. \square

We further modify \mathcal{A}_1 by exhaustively applying the following operation: if $p \in \text{Ind}(\mathcal{A}_1)$ is a non-root (i.e., p is a path with length > 1) and $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A}_p^-)} \models q_{\mathcal{T}}$,² where \mathcal{A}_p^- is obtained from \mathcal{A}_1 by removing the subtree rooted at p , then replace

² For the remainder of the proof, is it essential to consider all matches here instead of only the match π used to define \mathcal{A}_0 .

\mathcal{A}_1 with \mathcal{A}_p^- . Let \mathcal{A} be the Σ -ABox finally obtained. For $i \geq 0$, we say that \mathcal{A} has *outdegree at most i* if for every $p \in \text{Ind}(\mathcal{A}_1)$, the cardinality of the set

$$\{p' \in \text{Ind}(\mathcal{A}_1) \setminus \text{Ind}(\mathcal{A}_0) \mid \exists R : R(p, p') \in \mathcal{A}_1\}$$

is bounded by i . In words, every individual in \mathcal{A}_1 has at most i successors that are non-root nodes.

Lemma 8.

1. $\mathcal{I}_{f_{\mathcal{T}}^{m'}(\mathcal{A})} \models q_{\mathcal{T}}$, but $\mathcal{I}_{f_{\mathcal{T}}^{m'-1}(\mathcal{A})} \not\models q_{\mathcal{T}}$ for some $m' \geq m$;
2. \mathcal{A} has outdegree at most $|\mathcal{T}| + |q_{\mathcal{T}}|$.

Proof. For Point 1, it suffices to show that $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A})} \not\models q_{\mathcal{T}}$ and $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})} \models q_{\mathcal{T}}$. The former is a consequence of the facts that $\mathcal{I}_{f_{\mathcal{T}}^{m-1}(\mathcal{A}_1)} \not\models q_{\mathcal{T}}$, $\mathcal{A} \subseteq \mathcal{A}_1$, and $q_{\mathcal{T}}$ is a UCQ; the latter is immediate by construction of \mathcal{A} .

For Point 2, assume to the contrary of what is to be shown that there is a $p \in \text{Ind}(\mathcal{A})$ such that the cardinality of

$$\Gamma = \{p' \in \text{Ind}(\mathcal{A}) \setminus \text{Ind}(\mathcal{A}_0) \mid \exists R : R(p, p') \in \mathcal{A}\}$$

exceeds $|\mathcal{T}| + |q_{\mathcal{T}}|$. Fix a match π of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A})}$ and select all individual names in Γ that are in the range of π . For each $\exists r.A \sqsubseteq B \in \mathcal{T}$ such that $B(p) \in f_{\mathcal{T}}^{\infty}(\mathcal{A})$, select an individual name $p' \in \Gamma$ such that

1. $r(p, p') \in \mathcal{A}$ and $A(p') \in f_{\mathcal{T}}^j(\mathcal{A})$ for some $j < m$;
2. there is no p' that satisfies Point 1 for some smaller j .

Note that such a p' need not exist, in which case no node is selected for $\exists r.A \sqsubseteq B$. Since $|\Gamma| \geq |\mathcal{T}| + |q_{\mathcal{T}}|$, there is at least one element $p_0 \in \Gamma$ that is not selected. Consider the ABox $\mathcal{A}_{p_0}^-$ obtained from \mathcal{A} by dropping the subtree rooted at p_0 . Exploiting that \mathcal{T} is in normal form and using the definition of the function $f_{\mathcal{T}}(\cdot)$, it is not difficult to show that no deductions are lost at the individuals that remain in $\mathcal{A}_{p_0}^-$, i.e.,

Claim. For all $i \geq 0$, $p \in \mathcal{A}_{p_0}^-$, and concept names A , we have $A(p) \in f_{\mathcal{T}}^i(\mathcal{A}_{p_0}^-)$ iff $A(p) \in f_{\mathcal{T}}^i(\mathcal{A})$.

Consequently, π is still a match of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^m(\mathcal{A}_{p_0}^-)}$. This is a contradiction to the fact that the subtree rooted at p_0 has not been dropped during the construction of \mathcal{A} . \square

Since iterated applications of $f_{\mathcal{T}}$ can assign, in the worst case, every concept name from \mathcal{T} to every individual name, the following is immediate.

Lemma 9. *On every ABox \mathcal{B} with $|\text{Ind}(\mathcal{B})| = i$, we have $f_{\mathcal{T}}^{\ell}(\mathcal{B}) = f_{\mathcal{T}}^{\ell+1}(\mathcal{B})$ where $\ell = |\mathcal{T}| \cdot i$.*

We say that a path $p' \in \text{Ind}(\mathcal{A})$ is an *extension* of a path $p \in \text{Ind}(\mathcal{A})$ if $p' = pRa$ for some R and a . Set $d = |q_{\mathcal{T}}| + 3$. We claim that the ABox \mathcal{A} must contain a sequence

$$R_d(p_d, p_{d-1}), \dots, R_1(p_1, p_0)$$




Fig. 2. Gadget for the #P-hardness proof.

such that p_d is of length one and p_i is an extension of p_{i+1} for all $i < d$. Assume this is not the case. Then all paths in $\text{Ind}(\mathcal{A})$ are of length at most $d - 1$. Since \mathcal{A} is forest-shaped with at most $|q_{\mathcal{T}}|$ roots, Point 2 of Lemma 8 yields $|\text{Ind}(\mathcal{A})| \leq |q_{\mathcal{T}}| \cdot (|\mathcal{T}| + |q_{\mathcal{T}}|)^{d-1}$ and by Lemma 9 we have $f_{\mathcal{T}}^{\ell}(\mathcal{B}) = f_{\mathcal{T}}^{\ell+1}(\mathcal{B})$ where $\ell = |\mathcal{T}| \cdot |q_{\mathcal{T}}| \cdot (|\mathcal{T}| + |q_{\mathcal{T}}|)^{d-1}$. This is a contradiction to Point 1 of Lemma 8 and the fact that $m' > m > \ell$.

Lemma 10. *If at least one of $R_3(p_3, p_2)$, $R_2(p_2, p_1)$, $R_1(p_1, p_0)$ is dropped from \mathcal{A} resulting in an ABox \mathcal{A}' , then $\mathcal{A}', \mathcal{T} \not\models q$.*

Proof. Let $R_i(p_i, p_{i-1})$ be the assertion that was dropped from \mathcal{A} to obtain \mathcal{A}' and assume to the contrary of what is to be shown that $\mathcal{A}', \mathcal{T} \models q$. Thus $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A}')} \models q_{\mathcal{T}}$ by Theorem 11 and we can choose a concrete match π of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\mathcal{T}}^{\infty}(\mathcal{A}')}$. Since the length of p_{i-1} exceeds $|q_{\mathcal{T}}|$ and $q_{\mathcal{T}}$ is rooted (contains at least one individual name and is connected), none of p_{i-1}, \dots, p_0 can appear in the range of π and thus π is also a map from $\text{term}(q_{\mathcal{T}})$ to $\text{Ind}(\mathcal{A}_{p_{i-1}}^-)$. Trivially, we have $A(p) \in f_{\mathcal{T}}^{\infty}(\mathcal{A}')$ iff $A(p) \in f_{\mathcal{T}}^{\infty}(\mathcal{A}_{p_{i-1}}^-)$ for all $p \in \text{Ind}(\mathcal{A}_{p_{i-1}}^-)$ and concept names A . It follows that π is actually a match of $q_{\mathcal{T}}$ in $\mathcal{I}_{\mathcal{A}_{p_{i-1}}^-}$. This is a contradiction to the fact that the subtree rooted at p_{i-1} was not removed during the construction of \mathcal{A} . \square

We now prove Theorem 13 by a reduction of the problem of counting the number of satisfying assignments for a monotone bipartite DNF formula, which is known to be #P-hard. More specifically, input formulas are of the form

$$\psi = (x_{i_1} \wedge y_{j_1}) \vee \dots \vee (x_{i_k} \wedge y_{j_k})$$

where the set X of variables that occur on the left-hand side of a conjunction in ψ is disjoint from the set Y of variables that occur on the right-hand side of a conjunction in ψ .

For the reduction, let ψ be a formula as above, $X = \{x_1, \dots, x_{n_x}\}$, and $Y = \{y_1, \dots, y_{n_y}\}$. Define an ipABox $(\mathcal{A}_{\psi}, p_{\psi})$ by starting with the ABox \mathcal{A} constructed above and ‘multiplying’ the individuals a_1 and a_2 . In particular, introduce fresh individual names b_1, \dots, b_{n_x} (for a_2) and c_1, \dots, c_{n_y} (for a_1). Then add the following assertions to \mathcal{A}_{ψ} :

- add all assertions from \mathcal{A} that do not involve the individual names a_0, a_1, a_2 , or a_3 and assign probability 1 to all of them.

- add the following assertions with probability 1:

$$\begin{aligned}
& A(b_i) \text{ for all } A(a_2) \in \mathcal{A} \text{ and } 1 \leq i \leq n_x \\
& R(b_i, d) \text{ for all } R(a_2, d) \in \widehat{\mathcal{A}} \text{ and } 1 \leq i \leq n_x \\
& A(c_i) \text{ for all } A(a_1) \in \mathcal{A} \text{ and } 1 \leq i \leq n_y \\
& R(c_i, d) \text{ for all } R(a_1, d) \in \widehat{\mathcal{A}} \text{ and } 1 \leq i \leq n_y \\
& R_2(b_{i_\ell}, c_{j_\ell}) \text{ for } 1 \leq \ell \leq k;
\end{aligned}$$

where $\widehat{\mathcal{A}}$ is \mathcal{A} after dropping $R_3(a_3, a_2)$, $R_2(a_2, a_1)$, $R_1(a_1, a_0)$. See also Figure 2 for this part of the ABox.

- add the following assertions with probability 0.5,

$$\begin{aligned}
& R_3(a_3, b_i) \text{ for } 1 \leq i \leq n_x \\
& R_1(c_i, a_0) \text{ for } 1 \leq i \leq n_y.
\end{aligned}$$

We are interested in ABoxes $\mathcal{A}' \subseteq \mathcal{A}_\psi$ with $p_\psi(\mathcal{A}') > 0$. Each such ABox has probability $\frac{1}{2^{|\mathcal{X}|+|\mathcal{Y}|}}$ and corresponds to a truth assignment $\delta_{\mathcal{A}'}$ to the variables in $X \cup Y$: for $x_i \in X$, $\delta_{\mathcal{A}'}(x_i) = 1$ iff $R_3(a_3, b_i) \in \mathcal{A}'$ and for $y_i \in Y$, $\delta_{\mathcal{A}'}(y_i) = 1$ iff $R_1(c_i, a_0) \in \mathcal{A}'$. Let $\#\psi$ the number of truth assignments to the variables $X \cup Y$ that satisfy ψ . To complete the reduction, we show that $p(\mathcal{A}_\psi, \mathcal{T} \models q) = \frac{\#\psi}{2^{|\mathcal{X}|+|\mathcal{Y}|}}$. By what was said above, this is an immediate consequence of the following observation.

Lemma 11. *For all ABoxes $\mathcal{A}' \subseteq \mathcal{A}_\psi$ with $p_\psi(\mathcal{A}') > 0$, $\delta_{\mathcal{A}'} \models \psi$ iff $\mathcal{A}', \mathcal{T} \models q$.*

Proof. “if”. Let $\delta_{\mathcal{A}'} \not\models \psi$ and assume to the contrary of what is to be shown that $\mathcal{A}', \mathcal{T} \models q$. Since $\delta_{\mathcal{A}'} \not\models \psi$ and by construction of \mathcal{A}_ψ , there are no i, j such that $R_3(a_3, b_i), R_2(b_i, c_j), R_1(c_j, a_0) \in \mathcal{A}'$. Let \mathcal{A}'' be the restriction of \mathcal{A}' to those individual names that are reachable along role assertions from individual names that are roots, i.e., that occur already in \mathcal{A} and are paths of length one. Since q contains at least one individual name (which must be a root) and is connected, we have $\mathcal{A}'', \mathcal{T} \models q$. Let $\widehat{\mathcal{A}}$ be \mathcal{A} without $R_1(a_1, a_0)$ and let $h : \text{Ind}(\mathcal{A}'') \rightarrow \text{Ind}(\mathcal{A})$ be the identity on $\text{Ind}(\mathcal{A}'') \cap \text{Ind}(\mathcal{A})$ and such that $h(b_i) = a_2$ for all $b_i \in \text{Ind}(\mathcal{A}'')$ and $h(c_i) = a_1$ for all $c_i \in \text{Ind}(\mathcal{A}'')$. One can prove the following by induction on i .

Claim. For all $i \geq 0$, $p \in \text{Ind}(\mathcal{A}'')$, and $A \in \mathbf{N}_C$, $A(h(p)) \in f_{\mathcal{T}}^i(\mathcal{A}'')$ implies $A(p) \in f_{\widehat{\mathcal{A}}}^i(\widehat{\mathcal{A}})$.

Since $\mathcal{A}'', \mathcal{T} \models q$, by Lemma 6 there is a match π of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\widehat{\mathcal{A}}}}^\infty(\mathcal{A}'')$. By the above claim, π is also a match of $q_{\mathcal{T}}$ in $\mathcal{I}_{f_{\widehat{\mathcal{A}}}}^\infty(\widehat{\mathcal{A}})$, thus $\widehat{\mathcal{A}}, \mathcal{T} \models q$ by Lemma 6, in contradiction to Lemma 10.

“only if”. By construction of \mathcal{A}_ψ , $\delta_{\mathcal{A}'} \models \psi$ implies that, up to renaming of individual names that do not occur in q , we have $\mathcal{A} \subseteq \mathcal{A}'$. Since $\mathcal{A}, \mathcal{T} \models q$ by choice of \mathcal{A} , we must thus also have $\mathcal{A}', \mathcal{T} \models q$. \square

Theorem 7 (\mathcal{ELI} dichotomy). *Let q be a rooted Boolean CQ and \mathcal{T} an \mathcal{ELI} -TBox. Then q is in PTIME relative to \mathcal{T} or #P-hard relative to \mathcal{T} .*

Proof. If q is not FO-rewritable relative to \mathcal{T} , then it is #P-hard by Theorem 6. Using the semantics of CQs and of \mathcal{ELI} , it is not hard to show that FO-rewritings of CQs relative to \mathcal{ELI} -TBoxes are FO-formulas that are preserved under homomorphisms. By Rossman’s homomorphism preservation theorem on finite structures [29], we thus obtain that whenever there is an FO-rewriting for a CQ q and \mathcal{ELI} -TBox \mathcal{T} , then there is an FO-rewriting for q and \mathcal{T} that is a UCQ.³ Thus Theorem 1 and the PTIME vs. #P dichotomy for UCQs over tuple independent databases [10] immediately yields the dichotomy. \square

³ We are grateful to Meghyn Bienvenu for suggesting this argument.